



# AN APPLICATION OF MACHINE LEARNING TO COMPUTE THERMOCHEMISTRY OF REACTIVE FLOWS: A MIXTURE OF EXPERTS APPROACH

Reza Namdar<sup>1</sup>, Mohammad Norouzi<sup>1</sup>, Fathollah Varnik<sup>1</sup>

<sup>1</sup> Interdisciplinary Centre for Advanced Materials Simulation (ICAMS), Ruhr-University Bochum, Bochum, 44801, Germany. E-mail: fathollah.varnik@rub.de

## ABSTRACT

Accurate and efficient evaluation of chemical source terms is essential for high-fidelity simulations of reactive flows, particularly in large eddy simulation (LES) and direct numerical simulation (DNS). However, the stiffness and high dimensionality of detailed chemical kinetics render direct integration computationally expensive. To address this challenge, we propose a Mixture of Experts (MoE) neural network architecture that adaptively partitions the thermochemical space and models source terms using specialized subnetworks. These expert subnetworks are coordinated by a gating network that assigns input-dependent weights, allowing the model to capture complex, nonlinear reaction behavior across a wide range of conditions without manual zone definitions.

The MoE model is trained on data generated by detailed kinetic solvers and optimized using a mean squared error loss function. Once trained, it is integrated into our in house open source software OpenPhase Academic ([www.openphase.rub.de](http://www.openphase.rub.de)), allowing for fully self-contained combustion simulations without the need for external libraries such as CANTERA. Comparative simulations show that the MoE model reproduces reaction rates and thermochemical trends with high accuracy, closely matching the CANTERA reference results. Beyond accuracy, the MoE approach is more convenient comparing to other recommended methods for the implementation and is inherently compatible with GPU-based inference. This enables fast, on-the-fly thermochemistry evaluation, making the method highly suitable for large-scale reactive flow simulations.

**Keywords:** Thermochemistry, Neural Network, Mixture of Experts, Reactive Flows

## NOMENCLATURE

$N$	[-]	number of species
$n$	[-]	number of expert networks in the MoE model

$\mathcal{M}_k$	[-]	any property of species $k$
$\nu'$	[-]	molar stoichiometric coefficient of reactants
$\nu''$	[-]	molar stoichiometric coefficient of products
$\bar{W}$	[kg/mol]	mean molecular weight of the mixture
$\beta_j$	[-]	exponential temperature exponent for reaction $j$
$\dot{\omega}$	[W/m <sup>3</sup> ]	heat release rate
$\dot{\omega}_k$	[mol/m <sup>3</sup> · s]	net production rate of species $k$
$\underline{\mathbf{X}}$	[-]	input vector to the neural network
$\underline{\mathbf{Y}}$	[-]	predicted output vector from neural network
$\rho$	[kg/m <sup>3</sup> ]	density
$\sigma$	[-]	standard deviation of prediction error
$\theta_g$	[-]	trainable parameters of the gating network
$\theta_i$	[-]	trainable parameters of the $i$ -th expert network
$A_j$	[1/s]	pre-exponential factor for reaction $j$
$E_i(\underline{\mathbf{X}}; \theta_i)$	[-]	output of the $i$ -th expert network
$E_j$	[J/mol]	activation energy for reaction $j$
$g_i(\underline{\mathbf{X}})$	[-]	pre-activation output of gating network
$G_i(\underline{\mathbf{X}}; \theta_g)$	[-]	gating coefficient for $i$ -th expert
$K_{fj}$	[1/s]	rate constant of the forward reaction $j$
$K_{rj}$	[1/s]	rate constant of the reverse reaction $j$
$p$	[Pa]	thermodynamic pressure
$R$	[J/mol · K]	universal gas constant
$T$	[K]	temperature
$T_w$	[K]	wall temperature
$T_{inlet}$	[K]	inlet temperature
$W_k$	[kg/mol]	molecular weight of species $k$
$Y_k$	[-]	mass fraction of species $k$

$\mathcal{L}$        $[-]$       total loss function of the MoE model

### Subscripts and Superscripts

f      forward reaction  
inlet      inlet boundary  
j      index of reaction  
k      index of species  
r      reverse reaction  
w      wall

## 1. INTRODUCTION

Accurate and efficient modeling of chemical kinetics is essential for high-fidelity simulations of reactive flows, particularly in large eddy simulation (LES) and direct numerical simulation (DNS). These simulations depend on detailed chemical mechanisms to resolve flame structures and predict species evolution and heat release rates. However, the strong nonlinearity, stiffness, and high dimensionality of these mechanisms make direct integration of kinetic source terms prohibitively expensive in complex turbulent configurations. To reduce this cost, various surrogate modeling strategies have been explored, including mechanism reduction, manifold-based tabulation, and pre-trained neural networks [1, 2, 3, 4]. Among these, machine learning-based methods—particularly deep neural networks—have gained attention for their ability to approximate complex kinetics while significantly lowering computational expense.

One of the earliest applications of machine learning in combustion modeling employed dense feed-forward neural networks to emulate detailed chemical kinetics. These models are trained on high-fidelity datasets—typically generated from laminar flame simulations or precomputed manifolds—to learn mappings from input features such as species mass fractions and temperature to outputs like reaction rates or thermochemical source terms. Zhang et al. [5] demonstrated that deep neural networks can approximate complex reaction kinetics with high accuracy when trained on structured datasets and equipped with robust sampling strategies. Cheng et al. [4] further showed that such surrogates can be deployed in DNS to accelerate on-the-fly evaluation of chemical source terms, reducing overall simulation cost while preserving key flame characteristics. Despite their interpolation capabilities, these black-box models often lack physical consistency and may lose accuracy when extrapolating to regions of thermochemical space that are underrepresented in the training data [2].

To overcome the limitations of purely data-driven models, recent research has explored embedding physical laws into the training process through physics-informed neural networks (PINNs). Rather than relying solely on labeled data, PINNs incorporate domain knowledge—such as conservation of mass, element balance, and Arrhenius kinetics—into

the loss function, enabling the network to satisfy governing equations during optimization. Zhang et al. [6] introduced a constrained reaction-kinetics PINN (CRK-PINN), which embeds both element conservation and kinetic rate expressions to improve prediction accuracy in laminar flames. Wu et al. [7] extended this approach with FlamePINN-1D, which enables both forward and inverse solutions of 1D reacting flow problems while maintaining thermodynamic consistency and reducing data requirements. By embedding physical structure into the training process, PINNs offer improved extrapolation performance and generalizability, particularly in low-data or multiscale environments.

Moving beyond loss-function constraints, Ji and Deng [8] introduced the Chemical Reaction Neural Network (CRNN), a neural architecture explicitly designed to reflect the structure of chemical kinetics. Rather than imposing physical laws through the loss function, the CRNN encodes them directly into the network architecture. Each layer and connection corresponds to a chemical species or reaction pathway, allowing the model to learn both reaction mechanisms and kinetic parameters from time-series data of species concentrations. This physics-structured design enforces thermodynamic consistency and mass-action kinetics by construction, while also enhancing interpretability. However, the fixed structure of CRNNs may limit their flexibility when adapting to mechanisms with different topologies or under varying flow conditions.

While structured and physics-informed networks have improved the fidelity of combustion surrogates, another effective strategy involves decomposing the domain—spatially or in thermochemical space—into distinct regimes, each modeled by a separate neural network. Cheng et al. [4] implemented this by dividing the flame domain into four zones—burned gas, reaction zone, preheated zone, and unburned gas—based on local thermochemical structure. A separate neural network was trained for each zone, allowing specialization in the dynamics and scales relevant to each region. This clustering-based method reduced local approximation error and improved robustness in DNS. However, its reliance on manually defined thresholds tied to specific fuels and flame conditions limits generalization. Cheng et al. acknowledged this drawback and suggested that future work could benefit from automated or unsupervised clustering strategies adaptable to different reaction mechanisms or combustion regimes.

A natural progression of the zone-based clustering idea is the Mixture of Experts (MoE) framework, where both the partitioning of thermochemical space and the specialization of subnetworks are learned automatically during training. Owoyele et al. [9] applied this concept to surrogate modeling of combustion manifolds by training multiple expert networks and a gating function to emulate species mass fractions from precomputed flamelet data. Their method

eliminated the need for *a priori* clustering while preserving interpretability through data-driven specialization.

Building on this principle, the present work employs a MoE architecture to directly learn chemical source terms—specifically, reaction rate outputs from detailed kinetics—based on local thermochemical states. The proposed surrogate is designed for on-the-fly evaluation within reactive flow solvers, enabling accurate and efficient replacement of stiff chemical source term calculations without relying on predefined zonal boundaries. The primary motivation for employing a neural network model is to enhance the flexibility and portability of the simulation framework. Unlike traditional libraries such as Cantera, which are not compatible with GPU-based computation using CUDA, a neural network can be deployed entirely on the GPU. This eliminates the need for external software to evaluate thermochemical properties and enables high-performance simulations on GPU-accelerated platforms.

## 2. MATHEMATICAL MODEL

For a mixture containing  $N$  chemical species, the reactions occur according to [10]:

$$\sum_{k=1}^N \nu'_{kj} \mathcal{M}_k \rightleftharpoons \sum_{k=1}^N \nu''_{kj} \mathcal{M}_k, \quad (1)$$

where  $\mathcal{M}_k$  represents a property of species  $k$  involved in reaction  $j$ , and  $\nu'_{kj}$  and  $\nu''_{kj}$  are the molar stoichiometric coefficients of species  $k$  on the reactant and product sides, respectively.

Let  $Y_k$ ,  $\rho$ , and  $\dot{\omega}_k$  represent the mass fraction of species  $k$ , the mixture density, and the net production (or destruction) rate of species  $k$  due to chemical reactions. The governing equations are given by [10]:

$$\frac{\partial Y_k}{\partial t} = \frac{\dot{\omega}_k}{\rho}, \quad (2)$$

$$\dot{\omega}_k = W_k \sum_{j=1}^{N_r} \nu_{kj} \left( K_{fj} \prod_{k=1}^N [X_k]^{\nu'_{kj}} - K_{rj} \prod_{k=1}^N [X_k]^{\nu''_{kj}} \right), \quad (3)$$

$$\nu_{kj} = \nu''_{kj} - \nu'_{kj}, \quad (4)$$

$$[X_k] = \frac{\rho Y_k}{W_k}, \quad (5)$$

$$\rho = \frac{p \bar{W}}{RT}, \quad (6)$$

$$\frac{1}{\bar{W}} = \sum_{k=1}^N \frac{Y_k}{W_k}, \quad (7)$$

where  $W_k$  is the molecular weight of species  $k$ , and  $[X_k]$  is its molar concentration. The terms  $\bar{W}$ ,  $p$ , and  $T$  refer to the mean molecular weight, pressure, and temperature of the gas mixture, respectively. The universal gas constant is  $R = 8.314 \text{ J}/(\text{mol} \cdot \text{K})$ .



The reaction rate constants  $K_{fj}$  and  $K_{rj}$  for the forward and reverse directions of reaction  $j$  are typically modeled using the Arrhenius expression [10]:

$$K_j = A_j T^{\beta_j} \exp\left(-\frac{E_j}{RT}\right), \quad (8)$$

where  $A_j$ ,  $\beta_j$ , and  $E_j$  are empirical parameters determined from experimental measurements. All these constants are evaluated separately for forward and reverse reaction.

## 3. DATA GENERATION


In this study, data required for training the machine learning models were generated through detailed numerical simulations of premixed methane/air combustion, using the BFER [11] chemical kinetics mechanism. This mechanism involves six species:  $\text{O}_2$ ,  $\text{CH}_4$ ,  $\text{H}_2\text{O}$ ,  $\text{CO}$ ,  $\text{CO}_2$ ,  $\text{N}_2$ , participating in two global reaction steps,



For simplicity, the effect of pressure is neglected and fixed at 1 atm throughout all simulations.

To ensure comprehensive coverage of the thermochemical composition space encountered during practical simulations, a wall-flame interaction configuration was chosen as the test case. This configuration captures a wide range of flame behaviors and mixture conditions. The computational domain and boundary setup are shown in Figure 1. The key control parameters were the inlet temperature  $T_{\text{inlet}}$ , varied between 300 K and 800 K. The upper and the lower wall are kept at the same temperature,  $T_w$ . To cover a wide temperature range,  $T_w$  is varied from 300 K to 1200 K. These boundary conditions produce a wide variety of reactive flow states. To perform these simulations, OpenPhase Academic ([www.openphase.rub.de](http://www.openphase.rub.de)) was used, which was coupled with CANTERA to model the thermodynamic properties and kinetic information. The simulations were conducted using a hybrid numerical solver that combines the Lattice Boltzmann Method (LBM) with the Finite Difference Method (FDM). In this approach, the flow field is resolved using a thermal compressible Lattice Boltzmann model, while the temperature and species transport equations are solved using finite difference techniques. The underlying mathematical model describes a mixture comprising  $N$  species under low Mach number conditions, as detailed in [12]. From these simulations, the local flow field—specifically, species mass fractions and temperature—was sampled and recorded as input features. As the next step in generating the dataset, for the sake of higher accuracy, the corresponding output quantities, namely the net production rates of chemical species, were calculated via CANTERA. These quantities exhibit complex, nonlinear, and stiff behavior, primarily due to their strong temperature dependence following the Arrhenius law. The net production rates span several orders of magnitude, reflecting the multi-scale nature of chemical kinetics

and presenting a significant challenge for data-driven modeling.



**Figure 1.** Schematic view of the simulation domain used to generate training data. The setup features a premixed methane/air flow interacting with two planar parallel walls at constant and equal temperature. By varying the inlet temperature ( $T_{\text{inlet}}$ ) and wall temperature ( $T_w$ ), a wide range of thermochemical states are obtained within the gas phase, which are used to train the neural network models. We use the symmetry of the problem with respect to the middle plane and show only half of the domain for the fields of temperature (upper half) and the rate of heat release (lower half).

## 4. NEURAL NETWORK ARCHITECTURE

Inspired by the approach of Chi et al. [13], who manually partitioned the reactive flow domain into sub-regions, we employ a neural network architecture that learns such partitions automatically from data. Manual domain decomposition often relies on problem-specific heuristics that may not generalize across different chemical mechanisms. To address this, we adopt a Mixture of Experts (MoE) framework, which enables the network to learn distinct functional regimes in a data-driven and adaptive manner.

The MoE model consists of multiple specialized subnetworks, known as *experts*, and a separate *gating network* that dynamically assigns soft weights to each expert based on the input. This design enhances the model’s capacity to represent the complex, non-linear behavior characteristic of detailed chemical kinetics.

### 4.1. Problem Formulation

Let the input vector  $\underline{X} \in \mathbb{R}^d$  represent the local thermochemical state, consisting of the mass fractions of  $N$  chemical species and the temperature  $T$ :

$$\underline{X} = [Y_1, Y_2, \dots, Y_N, T, 1/T]^\top. \quad (9)$$

To better capture Arrhenius-type behavior, the reciprocal of temperature ( $1/T$ ) is included as an additional input. The output vector  $\underline{Y} \in \mathbb{R}^N$  contains the net production rates of the species, as computed by a detailed chemical kinetics solver.

In the Mixture of Experts framework, the target function is approximated using  $K$  expert networks. The final output is computed as a weighted sum of the expert predictions:

$$\underline{Y} = \sum_{i=1}^K G_i(\underline{X}; \theta_g) \cdot E_i(\underline{X}; \theta_i), \quad (10)$$

where  $E_i(\underline{X}; \theta_i)$  is the output of the  $i$ -th expert, and  $G_i(\underline{X}; \theta_g)$  is the corresponding gating weight computed by the gating network. These weights are normalized using a softmax function:

$$G_i(\underline{X}) = \frac{\exp(g_i(\underline{X}))}{\sum_{j=1}^K \exp(g_j(\underline{X}))}, \quad i = 1, \dots, K. \quad (11)$$

Each expert is encouraged to specialize in a particular thermochemical subdomain (e.g., high vs. low temperature, or fuel-lean vs. fuel-rich regions), while the gating network learns to interpolate between them smoothly.

### 4.2. Loss Function and Training Strategy

The model is trained end-to-end using supervised learning. The loss function is defined as the mean squared error (MSE) between the predicted and reference production rates:

$$\mathcal{L}(\theta_g, \theta_1, \dots, \theta_K) = \left\| \underline{Y}_{\text{true}} - \sum_{i=1}^K G_i(\underline{X}) E_i(\underline{X}) \right\|^2, \quad (12)$$

where  $\underline{Y}_{\text{true}}$  denotes the reference production rates from detailed kinetics. The model is optimized using the Adam optimizer. The dataset is split into training, validation, and test subsets.


### 4.3. Architectural Illustration

Figure 2 illustrates the overall architecture. The input vector is fed into both the gating and expert networks. The gating network computes soft weights for each expert, and the final output is the weighted sum of the expert outputs.

### 4.4. Advantages for Reactive Flow Modeling

The MoE framework offers several advantages for modeling reactive flows. First, it is dynamically adaptive and can respond to varying chemical regimes. Second, it promotes functional specialization: Each expert learns to make accurate predictions within a specific subdomain. This does not lead to a division of the physical space among the experts. Rather, each expert becomes a good predictor for a part of the thermochemical data space spanned by the two sets of the input and output vari-





**Figure 2.** Schematic of the Mixture of Experts (MoE) model (adapted from [14]). The input vector  $X$  is processed by a gating network that produces a soft distribution over  $n$  experts. Each expert yields an independent output vector, and the final output  $Y$  is computed as their weighted sum.

ables—specifically the distribution of net production rates as a function of species concentrations and temperature. Third, the use of softmax-based weighting ensures a smooth transition between expert outputs, preserving numerical stability. Finally, the modular architecture is scalable and can accommodate additional species or features with minimal structural modifications.

#### 4.5. Network Implementation Details


In our implementation, the input vector includes the mass fractions of five chemical species (excluding  $N_2$ ), the temperature  $T$ , and its reciprocal  $1/T$ , for a total of seven input features. The output vector contains the net production rates of the same five species.

Prior to training, both input and output values are normalized using Min-Max scaling. The MoE model includes five expert networks and a single gating network. Each expert comprises three hidden layers with 16 ReLU units per layer, followed by a linear output layer with five neurons. The gating network consists of one hidden layer with 32 ReLU units, followed by a softmax-activated output layer with five neurons to produce the expert weights. The entire dataset is divided into three parts. First, 20% of the data is set aside and not used until the final testing phase. The remaining 80% is then further split: 80% of it is used for training the model, and the remaining 20% is used for cross-validation.

The model is trained using the Adam optimizer and a mean squared error loss. Figure 3 compares the outputs of the trained model against reference values computed using CANTERA. As shown, the model accurately predicts source terms for unseen test data.

## 5. RESULTS AND DISCUSSION

In this section, we present the results of numerical simulations performed with the neural network-based thermochemistry model integrated into the OpenPhase solver. All simulations are carried out using OpenPhase, with the thermochemical source terms computed either from the trained neural network (NN) or using detailed chemistry from Cantera



**Figure 3.** Comparison of net production rates of species obtained from model prediction versus reference source terms computed with CANTERA for the test set.

for comparison. In all test cases, the combustion of a premixed methane/air mixture is simulated under stoichiometric conditions. The chemical kinetics are modeled using the reduced BFER mechanism, which offers a simplified yet sufficiently accurate depiction of essential combustion phenomena. Although the BFER mechanism is relatively compact, the advantages of employing a neural network (NN) framework become even more pronounced when applied to more complex kinetic mechanisms, where the computational overhead of traditional solvers such as Cantera increases significantly, while a trained NN can evaluate source terms with greater computational efficiency.

To evaluate the performance of the NN-based model, we consider two test cases: (i) a one-dimensional freely propagating flame, and (ii) a flame propagating through a packed bed of solid cylinders.

#### 5.1. One-Dimensional Flame Propagation


This test case simulates a freely propagating premixed flame in a one-dimensional domain. The initial setup consists of two regions: the left half contains a stoichiometric methane-air mixture, while the right half is filled with burnt combustion products.

An initial temperature perturbation is applied at the center to initiate ignition. The flame then propagates until a steady-state profile is achieved. The simulation parameters are summarized in Table 1.

**Table 1. Simulation parameters for the one-dimensional freely propagating flame case**

Parameter	Value
Inlet temperature	300 K
Pressure	1 atm
Equivalence ratio	1.0
Spatial resolution	$1 \times 10^{-5}$ m
Time step	$2 \times 10^{-8}$ s
Convective scheme	van Leer
Diffusion scheme	Central differencing
Fuel mixture	Methane/air
Chemical mechanism	BFER


Figure 4 shows the steady-state profiles of major species mass fractions obtained using the NN model and the detailed Cantera-based chemistry. Five key species are considered:  $\text{CH}_4$ ,  $\text{O}_2$ ,  $\text{CO}_2$ ,  $\text{H}_2\text{O}$ , and  $\text{CO}$ . The results from both models are in close agreement, indicating accurate prediction of species consumption and product formation.



**Figure 4. Comparison of species mass fraction profiles for the one-dimensional flame. Results from the NN model and Cantera are shown for  $\text{CH}_4$ ,  $\text{O}_2$ ,  $\text{CO}_2$ ,  $\text{H}_2\text{O}$ , and  $\text{CO}$ .**

In Figure 5, we compare the temperature profiles and heat release rate (HRR) for both models. The HRR peak location corresponds to the flame front and shows excellent agreement between the two approaches. The flame thickness and temperature rise are nearly identical, further confirming the NN

model's accuracy in predicting the flame structure.



**Figure 5. Comparison of temperature and heat release rate profiles for the one-dimensional flame. The peak of the HRR curve represents the flame front location.**

## 5.2. Flame Propagation in a Packed Bed


The second test case models the propagation of a premixed flame through a two-dimensional domain filled with solid cylinders, representing a packed bed. The computational domain is initially divided into two parts: the lower region contains a stoichiometric methane-air mixture, while the upper region contains combustion products. A small velocity is imposed at the inlet to allow the flame to propagate upstream toward the boundary. The simulation setup is summarized in Table 2.

**Table 2. Simulation parameters for the flame propagation in a packed bed**

Parameter	Value
Inlet temperature	300 K
Equivalence ratio	1.0
Inlet boundary	Fixed velocity
Outlet boundary	pressure outlet
Solid wall condition	No-slip, adiabatic
Spatial resolution	$2.5 \times 10^{-5}$ m
Time step	$2 \times 10^{-7}$ s
Convective scheme	van Leer
Diffusion scheme	Central differencing
Fuel mixture	Methane/air
Chemical mechanism	BFER

Figure 6 presents a sequence of snapshots showing the temporal evolution of the flame as it propagates through the packed bed. The flame front interacts with the solid structures, leading to local curvature changes that affect propagation dynamics. The NN-based and Cantera-based results are compared side-by-side to highlight differences in flame behavior and methane consumption.

Overall, both test cases demonstrate that the neural network model captures key features of combustion dynamics with good fidelity compared to detailed chemistry. Although the BFER mechanism is relatively simple, the NN-based approach offers sig-



**Figure 6.** Temporal progression of the flame front and methane mass fraction within the packed bed domain. The simulation employs a low inlet velocity at the lower boundary, with a pressure outlet condition specified at the upper boundary. Periodic boundary conditions are enforced along the lateral boundaries. Results are presented for the Cantera-based model (top row) and the neural network (NN) model (bottom row).

nificant potential for scaling to larger mechanisms and for deployment on GPU architectures — a direction we plan to explore in future work.

## 6. CONCLUSION

This work introduces a Mixture of Experts (MoE) neural network framework for efficient and accurate prediction of chemical source terms in reactive flow simulations. By leveraging multiple specialized expert networks and a gating network that adaptively partitions the thermochemical space, the proposed model captures the complex, nonlinear behavior of chemical kinetics across a wide range of combustion regimes. This data-driven approach eliminates the need for rigid zone definitions and improves flexibility compared to traditional surrogate models.

Trained on high-fidelity data from detailed chemical solvers such as Cantera, the MoE model demonstrates strong agreement with reference results. In addition to maintaining high accuracy, the architecture can significantly reduce computational overhead for large and stiff reaction mechanisms. Unlike traditional libraries, which are not GPU-compatible, the neural network-based implementation enables efficient inference on CUDA-capable hardware. This makes the method well-suited for large-scale, high-fidelity simulations, such as those encountered in DNS and LES of turbulent combustion.

Future work will focus on extending this approach to more complex mechanisms, incorporating advanced clustering strategies for expert assignment, and applying the framework to multi-dimensional flow configurations. Overall, the MoE architecture represents a promising step toward scalable, GPU-accelerated reactive flow modeling that preserves the fidelity of detailed chemistry while enabling real-time performance.

## ACKNOWLEDGEMENTS

Funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) - Project-ID 422037413 - TRR 287.

## REFERENCES

- [1] Li, K., Rahnama, P., Novella, R., and Somers, B., 2023, “Combining flamelet-generated manifold and machine learning models in simulation of a non-premixed diffusion flame”, *Energy and AI*, Vol. 14, p. 100266.
- [2] Readshaw, T., Franke, L. L., Jones, W., and Rigopoulos, S., 2023, “Simulation of turbulent premixed flames with machine learning - tabulated thermochemistry”, *Combustion and Flame*, Vol. 258, p. 113058.
- [3] Luo, L., Liu, Q., Sun, J., and Huang, Y., 2025, “Exploring surface reaction mechanism using a surface reaction neural network framework”, *Chemical Engineering Science*, Vol. 306, p. 121307.
- [4] Chi, C., Janiga, G., and Thévenin, D., 2021, “On-the-fly artificial neural network for chemical kinetics in direct numerical simulations of premixed combustion”, *Combustion and Flame*, Vol. 226, pp. 467–477.
- [5] Zhang, T., Yi, Y., Xu, Y., Chen, Z. X., Zhang, Y., E, W., and Xu, Z.-Q. J., 2022, “A multi-scale sampling method for accurate and robust deep neural network to predict combustion chemical kinetics”, *Combustion and Flame*, Vol. 245, p. 112319.
- [6] Zhang, S., Zhang, C., and Wang, B., 2024, “CRK-PINN: A physics-informed neural network for solving combustion reaction kinetics ordinary differential equations”, *Combustion and Flame*, Vol. 269, p. 113647.
- [7] Wu, J., Zhang, S., Wu, Y., Zhang, G., Li, X., and Zhang, H., 2025, “FlamePINN-1D: Physics-informed neural networks to solve forward and inverse problems of 1D laminar flames”, *Combustion and Flame*, Vol. 273, p. 113964.
- [8] Ji, W., and Deng, S., 2021, “Autonomous Discovery of Unknown Reaction Pathways from Data by Chemical Reaction Neural Network”,

*The Journal of Physical Chemistry A*, Vol. 125 (9), pp. 1990–1999.

- [9] Owoyele, O., Kundu, P., and Pal, P., 2021, “Efficient bifurcation and tabulation of multi-dimensional combustion manifolds using deep mixture of experts: An a priori study”, *Combustion and Flame*, Vol. 226, pp. 203–215.
- [10] Poinso, T., and Veynante, D., 2005, *Theoretical and numerical combustion*, RT Edwards, Inc.
- [11] Franzelli, B. G., 2011, “Impact of the chemical description on direct numerical simulations and large eddy simulations of turbulent combustion in industrial aero-engines”, Ph.D. thesis, INPT.
- [12] Hosseini, S. A., Safari, H., Darabiha, N., Thévenin, D., and Krafczyk, M., 2019, “Hybrid lattice Boltzmann-finite difference model for low Mach number combustion simulation”, *Combustion and Flame*, Vol. 209, pp. 394–404.
- [13] Chi, C., Janiga, G., and Thévenin, D., 2021, “On-the-fly artificial neural network for chemical kinetics in direct numerical simulations of premixed combustion”, *Combustion and Flame*, Vol. 226, pp. 467–477.
- [14] Alazizi, A., Habrard, A., Jacquenet, F., He-Guelton, L., and Oblé, F., 2020, “Dual sequential variational autoencoders for fraud detection”, *Advances in Intelligent Data Analysis XVIII: 18th International Symposium on Intelligent Data Analysis, IDA 2020, Konstanz, Germany, April 27–29, 2020, Proceedings 18*, Springer, pp. 14–26.