State prediction of freely propagating flames using Fourier neural operators

By F. Rios Tascon, R. F. Johnson†, D. D. Ortiz, P. J. Schmid‡ AND B. J McKeon

The simulation of high-speed chemically reacting flows is a computationally expensive task because of the number of species that need to be accounted for to obtain the complete state of the flow. Reduced order models (ROMs) are a tool that can potentially alleviate some of this computational burden. This brief explores the use of Fourier neural operators (FNOs) as a ROM for the case of a freely propagating, adiabatic, premixed flame. Key thermodynamics, and some major chemical species are used as inputs in an effort to predict the rest of the species that would give a complete set of parameters for the flow.

1. Introduction

The accurate simulation of chemically reactive flows is still a large challenge in computational fluid dynamics (CFD). High-fidelity simulations of combustion processes are computationally expensive due to the stiffness and high dimensionality of the governing equations (Poinsot & Veynante 2005). Due to the number of species involved in these types of simulations, as well as the fine spatial and temporal resolution required to capture the gradients near areas such as flame fronts, even simple configurations can be prohibitively expensive.

An approach that promises to alleviate some of this computational burden is the use of reduced order models (ROMs). ROMs have been used successfully to accelerate chemistry solvers (Lu & Law 2009). However, a problem with many traditional ROM techniques is that they rely on assumptions that may not always transition to macroscopic objectives which can operate outside of the confines in which the ROM was built. Recent advances in machine learning have opened up the possibility for the development of new ROMs. In particular, Fourier neural operators (FNOs) (Li *et al.* 2021) have been demonstrated to learn mappings between infinite-dimensional function spaces with high accuracy, and thus it is reasonable to explore their deployment for state completion in combustion calculations.

This work presents an FNO-based ROM for one-dimensional, freely propagating, adiabatic premixed hydrogen flames. This ROM was trained on over 15,000 flow cases generated using Cantera (Goodwin *et al.* 2025), and the detailed hydrogen sub-mechanism was extracted from the FFCM-2 chemical model (Zhang *et al.* 2023). Using a two-stage FNO architecture, it is shown that the full thermochemical state of the flame can be reconstructed from a reduced set of input variables. The model is trained and validated on a wide range of pressures, temperatures, and equivalence ratios. The performance of this FNO-based ROM is then evaluated using two different metrics. Lastly, a brief discussion on the modular and interpretable nature of this FNO architecture is included.

† U.S. Naval Research Laboratory
‡ King Abdullah University of Science and Technology

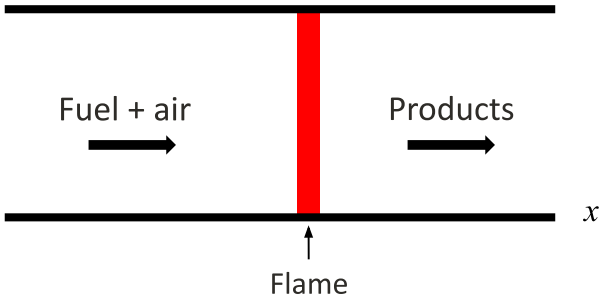

FIGURE 1. Freely propagating, adiabatic, premixed flame setup.

Parameter	Parameter range
T	300–800 K
p	0.5–3 atm
ϕ	0.3–1.3

TABLE 1. Ranges of temperatures (T), pressures (p), and equivalence ratios (ϕ) used as parameters in Cantera to generate more than 15,000 flow cases.

2. Approach

The work presented here focuses on a freely propagating, adiabatic, premixed flame. A diagram of this flow is shown in Figure 1. In this case, hydrogen (H_2) is used as fuel premixed in air (O_2 and N_2). The data pertaining to this flow were generated using Cantera (Goodwin *et al.* 2025), an open-source chemical kinetics software. More than 15,000 flow cases were generated over a range of temperatures (T), pressures (p), and equivalence ratios (ϕ). The ranges for these parameters are outlined in Table 1.

The flow cases generated using Cantera give a spatial distribution of nine different mass fractions, as well as a temperature profile, along the axis of the flame. These ten parameters then define the complete state of the flow in each instance. The Foundational Fuel Chemistry Model Version 2.0 (FFCM-2) reaction model (Zhang *et al.* 2023) was implemented in the Cantera simulation. In this particular case, 25 reaction equations were used to model the combustion process.

The equivalence ratio, pressure, and temperature parameters determine not only the shape of the mass fraction profiles, but also the x location of the flame itself. Figure 2 shows the distribution of the flame location as the different governing parameters are varied. The flame is located in the vicinity of $x = 0.3$, with case-to-case variation. Knowledge of the location of the flame was used to preprocess the data, as explained in the next section.

An FNO-based model was implemented to predict seven of the nine species mass fractions present in this flow, using the other two species and temperature as inputs to the model. FNOs implement Fourier transforms to learn spatial patterns in the data (Li *et al.* 2021). Specifically, they consist of a non-linear encoder and decoder, along with (normally) a series of Fourier layers with non-linear activation functions that learn an operator on the most energetic Fourier modes for each layer. In our case, only one linear

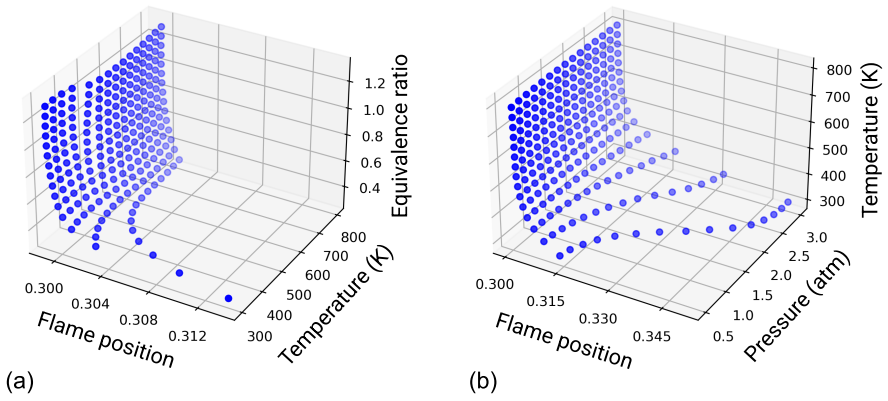

FIGURE 2. (a) Distribution of flame positions with varying temperature and equivalence ratio for a fixed pressure ($p = 0.5$ atm), and (b) the distribution of flame positions with varying temperature and pressure for a fixed equivalence ratio ($\phi = 0.3$).

Fourier layer was used to add some modularization and interpretability to our model compared to traditional FNOs.

3. Data preprocessing

As shown in Figure 3, a flame front can be divided into three distinct regions (Peters 2010). In the first region, the preheat zone, the fuel and oxidizer have not yet combusted, but the temperature gradually starts to increase. The second region, the reaction zone, is the section of the flame in which most of the combustion process takes place. This zone is characterized by high gradients in temperatures and mass fraction, which requires the finest grid spacing, as shown in Figure 3(b). This figure also shows the x location where the gas reaches 10% ($x_{10\%}$) and 90% ($x_{90\%}$) of the temperature rise from its unburned to its burned state, which shows how most of the temperature rise happens within a very small region near the reaction zone. In the final region, the post-flame zone, most of the products have been formed, and there are no high gradients present.

FNOs use a Fourier transform as a tool to learn the most prominent features of the data. Because of this, they can struggle with high gradients in the data, where the Gibbs phenomenon might occur. To address this issue, the one-dimensional flame front cases generated using Cantera were preprocessed before training the FNO. Knowledge of the different regions of the flame was used to regrid the data in order to add resolution to the average reaction zone location. The distribution of grid spacing over the entire x-axis is shown in Figure 3(a). The spacing used in this grid was generated using a Gaussian distribution function with mean at the average flame location ($x \approx 0.3$), and standard deviation of $\sigma = 0.03$. Similar results to the ones presented in this brief could be obtained with the use of other grids, as long as their spacing is fine enough near the reaction zone to smooth out the sharp gradients.

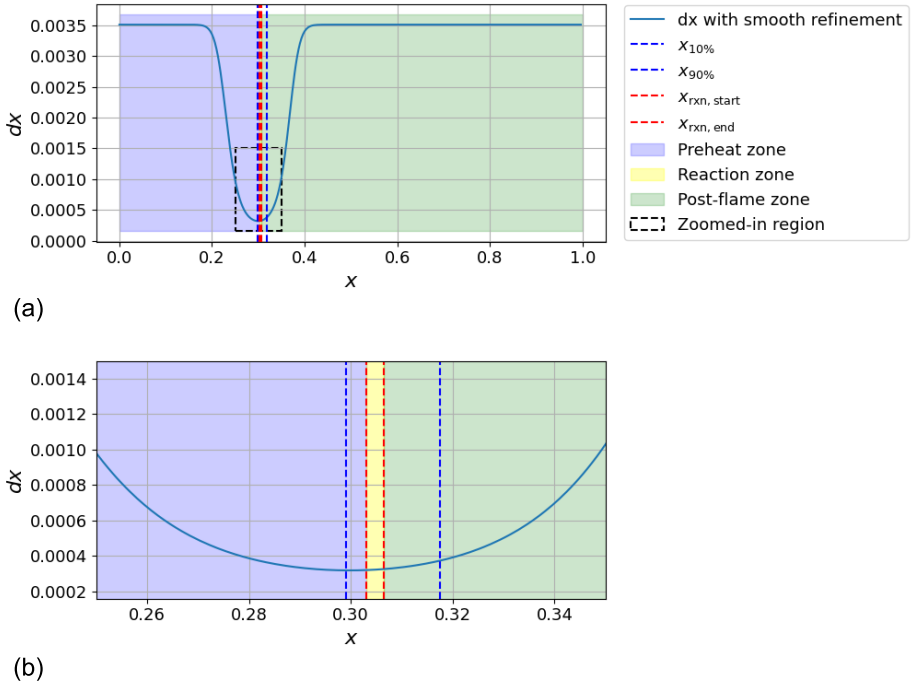

FIGURE 3. (a) Grid spacing over the entire flame and (b) grid spacing in the zoomed-in region near the reaction zone.

4. State-completing FNO architecture

The two-stage FNO architecture shown in Figure 4 was implemented to demonstrate state completion using only temperature and the H_2 and O_2 mass fractions on the input side. The first FNO, shown in Figure 4(a), takes temperature and the mass fractions for H_2 and O_2 (Y_{H_2} and Y_{O_2}), and predicts the mass fractions of the three main species present in the post-flame zone (Y_{H_2O}, Y_{OH}, and Y_O). In the second FNO, shown in Figure 4(b), temperature and the mass fractions for H_2 and O_2 are once again used as inputs, but now the set of input variables are expanded by adding the predicted mass fraction for H_2O from the first FNO as one of the inputs. Using this expanded set of inputs, the mass fraction for minor species in the flow, those that are present mainly near the reaction zone (Y_{N_2}, Y_H, Y_{HO_2}, and $Y_{H_2O_2}$), are predicted.

Results were obtained by training the FNOs using 10,000 training flow cases, and 3,000 validation flow cases. An additional 2,000 cases were left as test data. During the training process, an early stopping criterion was implemented to avoid overfitting. In this case, the training of the model was stopped once the validation error converged, or more specifically, after five consecutive epochs of no significant improvement in the validation error. This criterion was met after approximately 25 epochs.

Each of the two FNOs consist of a non-linear encoder (P), a non-linear decoder (Q), and a single linear Fourier layer. Inside the top section of the Fourier layer of each FNO,

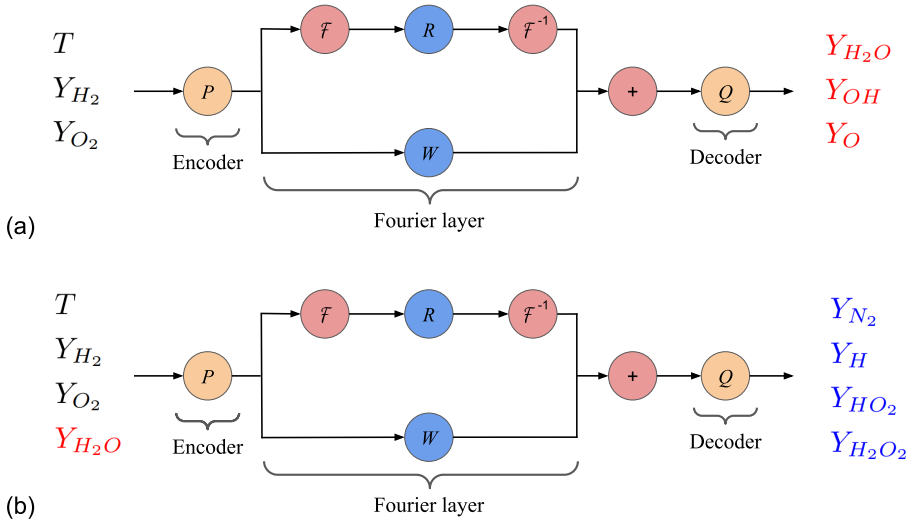

FIGURE 4. (a) First FNO mapping, which predicts the mass fractions for the main species, and (b) econd FNO mapping, which takes the expanded set of input variables to predict the mass fractions for the minor species.

shown in Figure 4, an operator R is learned for the most energetic Fourier modes of the input; the bottom section (W) acts as a bias (Li *et al.* 2021). Having a single linear Fourier layer in each FNO not only helps ensure that the networks are not unnecessarily deep, which in turn helps avoid overfitting, but also opens up the possibility of using this FNO-based ROM as an input-output model to learn functional relationships between different subprocesses in the data.

One of the main advantages of a prediction in two stages over a single stage is the potential for modularization during the training and prediction process. It also allows for a slightly more interpretable way of understanding the mapping between the inputs and outputs. Neither of these are explored in this brief, but they are discussed in Section 6.

For this architecture, mean squared error (MSE) was implemented as part of the training procedure. This loss is defined as follows

$$\text{MSE} = \frac{1}{n}\sum_{i=1}^{n}(Y_{\text{true},i} - Y_{\text{pred},i})^2, \tag{4.1}$$

where Y_{true} and Y_{pred} are the true and predicted mass fractions respectively, and n is the number of grid points.

5. Results

Figure 5(a-c) shows the prediction of the mass fractions for H_2O, OH, and O from the first FNO mapping for a single representative test case in the test dataset. The FNO is able to correctly predict the location of the flame, and it somewhat accurately predicts the shape of each profile.

Figure 5(d-g) then shows the predictions for the remaining four species that complete the state of the flow. The FNO is able to predict the location of the flame, and although

there appears to be some more substantial error compared to the first mapping, the magnitude of each of these species is very small, so it has a small effect on the overall error for the predicted complete set of mass fractions. As mentioned in Section 3, FNOs generally struggle with high gradients in the data, such as the one seen in Figure 5(d) for the N_2 mass fraction. It is important to note that N_2 being inert, its mass fraction should be constant throughout the entire domain, but in this case, it is not constant due to numerical errors for which Cantera accounts by changing the mass fraction of N_2 to ensure that the sum of all mass fractions still equals one. Since this discontinuity is present in the training data, it is therefore also learned by the FNO.

The performance of this FNO-based ROM can be characterized by comparing the sum of all the species mass fractions (the two used as inputs plus the seven predicted) with the true value of one. That is, by the definition of the mass fraction, it is expected that

$$\sum_i Y_i = 1, \tag{5.1}$$

where Y_i is the mass fraction of each species. The average computed and predicted sum of mass fractions over 31 test flow cases is shown in Figure 6. This figure shows that the model is on average able to correctly ensure that the sum of mass fractions at any spatial location is close to one. As expected, most of the error occurs near the reaction zone of the flame, but on average the error is still less than 1%. Even though Equation (5.1) was not directly enforced in the model, it is implicitly being learned. The possibility of enforcing physics constraints like this during the training of the model is discussed in the next section.

To quantify the error of each of the predicted mass fractions, the weighted mean squared error (WMSE) was used. In this case, the weights correspond to the grid spacing (dx_i), as shown in Figure 3, so the WMSE is then defined as

$$\text{WMSE} = \frac{1}{n}\frac{\sum_{i=1}^{n} dx_i (Y_{\text{true},i} - Y_{\text{pred},i})^2}{\sum_{i=1}^{n} dx_i}, \tag{5.2}$$

where Y_{true} and Y_{pred} are the true and predicted mass fractions, respectively, and n is the number of grid points. The average WMSE for each of the seven predicted mass fractions over 31 test flow cases are shown in Table 2. Since the WMSE on its own can be hard to interpret, the maximum error at any x location for each of the species was calculated as

$$\text{Max error} = \max_i |Y_{\text{true},i} - Y_{\text{pred},i}|, \tag{5.3}$$

The average of the maximum error over 31 test flow cases is also shown in Table 2. The variance of the WMSE and of the maximum error for the 31 test cases were also calculated as follows

$$\text{Variance of WMSE} = \frac{1}{m}\sum_{j=1}^{m}(\text{WMSE}_j - \overline{\text{WMSE}})^2, \tag{5.4}$$

$$\text{Variance of max error} = \frac{1}{m}\sum_{j=1}^{m}(\text{max error}_j - \overline{\text{max error}})^2, \tag{5.5}$$

where m is the number of test cases (in this case $m = 31$), WMSE_j and max error_j are the errors for the j-th test case, and $\overline{\text{WMSE}}$ and $\overline{\text{max error}}$ are the mean errors for all test cases. These variances are in Table 2. These two metrics and their variances show

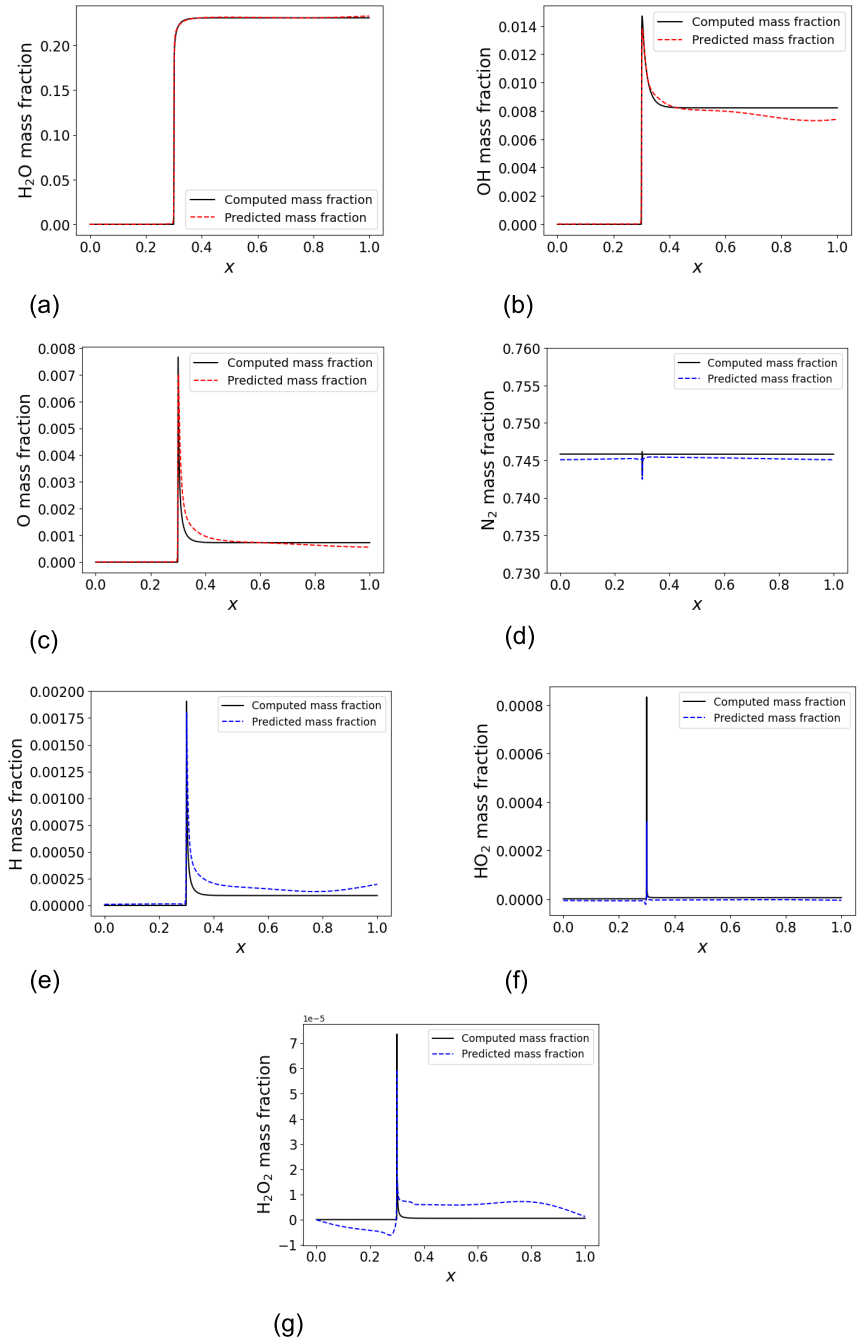

FIGURE 5. Mass fraction predictions for a single representative test case using the first FNO mapping (red dashed lines) for (a) H_2O, (b) OH, and (c) O and using the second FNO mapping (blue dashed lines) for (d) N_2, (e) H, (f) HO_2, and (g) H_2O_2.

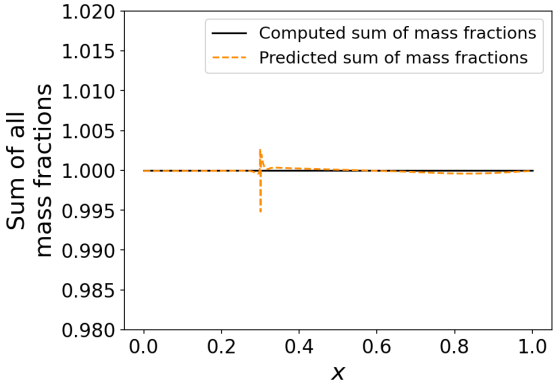

FIGURE 6. Average of the sum of the computed and predicted mass fractions for 31 test flow cases.

Species mass fraction	Average WMSE	Variance of WMSE	Average max error	Variance of max error
H_2O	1.88×10^{-9}	8.51×10^{-18}	1.45×10^{-2}	4.69×10^{-8}
OH	3.56×10^{-9}	4.67×10^{-17}	3.84×10^{-3}	5.87×10^{-6}
O	4.91×10^{-10}	7.27×10^{-19}	2.49×10^{-3}	1.25×10^{-6}
N_2	4.02×10^{-10}	2.32×10^{-19}	1.76×10^{-3}	3.55×10^{-7}
H	1.16×10^{-11}	2.39×10^{-22}	4.19×10^{-4}	7.03×10^{-8}
HO_2	4.46×10^{-13}	8.75×10^{-26}	4.06×10^{-4}	2.21×10^{-8}
H_2O_2	7.59×10^{-14}	1.62×10^{-26}	7.25×10^{-5}	1.46×10^{-9}

TABLE 2. Average weighted mean square error (WMSE) and average maximum error with their variance for each of the seven predicted mass fractions over 31 test flow cases. The average WMSE and maximum error were computed by taking the average over the 31 test cases of their values calculated using Eqs. (5.2) and (5.3) respectively. Their variance was then computed using Eqs. (5.4) and (5.5), respectively.

that the predicted mass fractions are encouragingly close to the mass fractions computed using Cantera.

6. Discussion and future work

The proof of concept results in the previous section show that FNOs offer significant potential to serve as ROMs for chemically reacting flows. The presented architecture accurately predicts full thermochemical states from a reduced set of inputs across a wide range of conditions, while maintaining low prediction error even near the reaction zone. It is important to note that other types of network architectures could have a comparable (or better prediction) performance than what has been shown here, but this particular FNO-based ROM has several key advantages, particularly its modularization and interpretability. Also, similar to other machine learning-based ROMs, this FNO model allows for the integration of physics into the model.

6.1. *Modularization*

A core design choice in this work is the use of a two-stage FNO architecture, where the prediction of major species is decoupled from that of minor species. This modular architecture has several advantages:

- Targeted learning: By partitioning the predictive problem into smaller subproblems, each FNO can be trained to be tailored toward the different subsets of species with distinct physical roles. The first FNO focuses on major products that dominate the energy release and overall structure of the flame (i.e., H_2O, OH, and O), while the second targets more transient or intermediate species (e.g., HO_2 and H_2O_2). This separation allows for the implementation of a different loss function that targets the physics or numerical characteristics of each group of species.
- Robustness: This type of network permits separation of the learning of major and minor species, adding robustness to the learning process.

Future work will explore more advanced connections between the different parts of the FNO, for example incorporating hierarchical feedback mechanisms where the outputs of one FNO inform the inputs of another, or training them jointly using multi-task loss functions.

6.2. *Interpretability*

The network structure also leads to relative interpretability of the model. This is particularly because of the single linear Fourier layer, which allows for the study of the latent space. The following aspects aid the interpretability of the model:

- Fourier layer analysis: Having a single linear Fourier layer allows for the potential insights regarding the mapping from the inputs to the outputs, but the associated physical space relationships enabling prediction of the minor species will require an additional interpretation step.
- Input-output mapping: The modular architecture and reduced input and output spaces enable examination of direct functional relationships between specific inputs and outputs. For example, a study could be done into how the variations of H_2 or O_2 mass fractions shifts the predicted OH or HO_2 profiles. This would give insights into the sensitivities and functional relationship between subprocesses.

As a future direction, an input-output analysis could be performed to determine the sensitivity of each species to other species, and to gain an understanding of the functional relationship between different subprocesses in the flow.

6.3. *Integrating physics into the model*

The model presented in this brief was purely data-driven, but as shown in Section 5, it was able to implicitly learn physical constraints such as conservation of mass (i.e., sum of mass fractions $\sum_i Y_i = 1$), even though this constraint was not explicitly enforced during training. A next step would be to embed physical principles into the model during the training process. This could further improve model accuracy and generalizability.

One common approach for integrating physics into the model is by utilizing physics-informed loss functions. In this case, a constraint such as mass conservation would be included as an additional term in the loss function. The model would then penalize any outputs that do not satisfy $\sum_i Y_i = 1$. Future work will include the implementation of these types of physics constraints into the model.

6.4. *Extending the method to other data*

The method outlined here was developed for the study of premixed flames. However, because of the modular nature of this type of model, with some modifications, this method could be applied to a variety of combustion scenarios. To use this architecture for a multi-component flame, which involve the combustion of multiple fuel types, the FNO-based ROM would need to be modified in one of two ways: either increasing the input and output size of each of the two FNOs, or having a third FNO (or more) that learns the extra species. For non-adiabatic flames additional inputs to the model could be added such as temperature gradients, or alternatively, boundary condition information could be incorporated into the loss function. Finally, for unsteady flames, since the location of the flame front would change with time, there would have to be some adaptive gridding done to smooth out the steep gradients.

This FNO-based ROM then offers a powerful and scalable approach to combustion modeling for a variety of flame cases. If incorporated as a ROM into CFD codes, it could enable the rapid prediction of species, while also being suited for use as a state-completion model in experimental settings where only a portion of all species are measured.

Acknowledgments

The support of the Office of Naval Research under grant number N00014-23-1-2646 is gratefully acknowledged.

REFERENCES

Goodwin, D., Moffat, H., Speth, R. & Weber, B. 2025 Cantera: An object-oriented software toolkit for chemical kinetics, thermodynamics, and transport processes [software]. version 3.2.0. *Cantera Developers* .

Li, Z., Kovachki, N., Azizzadenesheli, K., Liu, B., Bhattacharya, K., Stuart, A. & Anandkumar, A. 2021 Fourier neural operator for parametric partial differential equations. *arXiv:2010.08895* .

Lu, T. & Law, C. 2009 Toward accommodating realistic fuel chemistry in large-scale computations. *Prog. Energy and Combust. Sci.* **35**, 192–215.

Peters, N. 2010 Combustion theory - reference notes. *Combustion Energy Frontier Research Center* .

Poinsot, T. & Veynante, D. 2005 Theoretical and numerical combustion, second edition. *R.T. Edwards* p. 177.

Zhang, Y., Dong, W., Vandewalle, L., Xu, R., Smith, G. & Wang, H. 2023 Foundational fuel chemistry model version 2.0 (ffcm-2). *Stanford University* .