

RESEARCH ARTICLE

Physics-constrained machine learning for reduced composition space chemical kinetics

Anuj Kumar and Tarek Echekki 

Department of Mechanical and Aerospace Engineering, North Carolina State University, Raleigh, NC, USA

Corresponding author: Tarek Echekki; Email: techekk@ncsu.edu

Received: 04 June 2024; **Revised:** 26 March 2025; **Accepted:** 31 March 2025

Keywords: chemistry acceleration; deep operator networks; physics-informed machine learning; reduced order modeling; surrogate modeling

Abstract

Modeling detailed chemical kinetics is a primary challenge in combustion simulations. We present a novel framework to enforce physical constraints, specifically total mass and elemental conservation, during the reaction of ML models' training for the reduced composition space chemical kinetics of large chemical mechanisms in combustion. In these models, the transport equations for a subset of representative species are solved with the ML approaches, while the remaining nonrepresentative species are "recovered" with a separate artificial neural network trained on data. Given the strong correlation between full and reduced solution vectors, our method utilizes a small neural network to establish an accurate and physically consistent mapping. By leveraging this mapping, we enforce physical constraints in the training process of the ML model for reduced composition space chemical kinetics. The framework is demonstrated here for methane, CH₄, and oxidation. The resulting solution vectors from our deep operator networks (DeepONet)-based approach are accurate and align more consistently with physical laws.

Impact statement

The integration of complex fuel chemistry in practical computational fluid dynamics (CFD) applications requires extensive computational resources. We integrate two machine-learning frameworks to accelerate chemistry integration, reduce the dimensionality of the solution vector, reconstruct the full chemical system, and enforce mass and element conservation. The first framework, React-DeepONet, advances the solution of representative species by mapping it across small time increments. The second framework, Correlation Net, reconstructs the remaining species and implements mass and element conservation to evaluate training loss. The integrated framework reduces the computational cost and memory requirements associated with advancing a full chemical system. At the same time, Correlation Net emulates the output of a full chemical system solution with mass and element conservation constraints.

1. Introduction

The detailed modeling of chemical kinetics is the principal bottleneck in the combustion simulations. Recently, several efforts have been made to learn the chemical kinetics and accelerate its integration in a simulation using machine learning (ML) tools (Blasco et al., 1998, 2000; Haghshenas et al., 2021; Ji and Deng, 2021; Ranade et al., 2019). However, due to the high dimensionality of the chemical systems for

complex fuels, the ML models have predominantly been restricted to simpler fuels and mechanisms. For larger and more complex reaction mechanisms that have a high-dimensional solution vector of thermo-chemical scalars, chemical kinetics may be learned in the reduced composition space (Dikeman et al., 2022; Kumar and Echekki, 2024, either through dimensionality reduction techniques by advancing solutions of the chemical system in latent space (see, for example, (Goswami et al., 2024; Kumar and Echekki, 2024) or by transporting a subset of the chemical system through representative species (Kumar and Echekki, 2024).

While ML models have demonstrated significant potential for capturing chemical kinetics in the reduced composition space, the temporal evolution of the reduced solution vector may not conform to certain physical constraints, like total mass and elemental conservation. These physical constraints are typically essential for CFD solvers, potentially hindering the integration of ML models with such solvers.

Recent efforts have explored incorporating domain-specific physical principles or laws into neural network training by adjusting the loss function (Raissi et al., 2019), allowing the models to rely on data and established physical knowledge. An obvious step of incorporating the physical constraint of total mass and elemental conservation in the ML training process for reduced composition space chemical kinetics is hindered by the unavailability of the full solution vector representing all species in a chemical mechanism. Enforcing these constraints during training often requires the recovery of the full solution vector.

In this article, we introduce a framework that enables the implementation of a chemistry acceleration tool based on predicting a subset of species within the chemical system. This subset is referred to as representative species from which other nonrepresentative species can be recovered. Using training data, this framework recovers nonrepresentative species using an artificial neural network (ANN) called Correlation Net and implements mass and element conservation within the loss evaluation.

This framework is readily adaptable to any ML model aiming to capture the chemical kinetics of complex combustion fuels with a subset of the species. However, in this study, we integrate it with the React-DeepONet of Kumar and Echekki (2024). The React-DeepONet architecture is based on deep operator networks (DeepONets) by Lu et al. (2021). React-DeepONet is designed to map a solution vector of a subset of the thermo-chemical scalars, temperature, and representative species at a prescribed time onto a solution at a later small time increment. Representative species are selected because of their role as markers for the evolution of fuel oxidation and their strong correlation with the remaining (nonrepresentative species) (Pope, 2013). Similar approaches using DeepONets, albeit with the integration of the full chemical systems, have been proposed by Venturi and Casey (2023) and Goswami et al. (2024). As in these two approaches, React-DeepONet is trained on simulation data. Its implementation with a reduced set of chemical systems is designed to overcome the challenges of integrating a full thermo-chemical system with stiff ODEs for large and complex fuels. The React-DeepONet-based scheme results in up to three orders of magnitude accelerations in the chemistry prediction of complex fuels (Kumar and Echekki, 2024).

In this study, we first establish a physically consistent mapping from the reduced set of representative species to the full solution vector. Subsequently, we incorporate the physical constraints of total mass and elemental conservation into the training of React-DeepONet, which learns chemical kinetics in a reduced composition space. As a result, the solution vector produced by React-DeepONet achieves greater physical consistency, retaining its predictive accuracy.

2. Methodology

2.1. Chemical kinetics identification through a deep operator network

This section discusses the proposed framework for the physics-informed enforcement of mass and element conservation. As stated earlier, we integrate the framework with our recently developed React-DeepONet framework for chemistry acceleration. The framework is designed to predict the solution of a

subset of the thermo-chemical scalars' vectors between two prescribed time increments. Therefore, the framework is designed to advance the chemical source operator that can be coupled with transport in a CFD simulation. We illustrate the framework using data for a 0D, homogeneous mixture with prescribed pressure, composition (through an equivalence ratio), and initial temperature. The corresponding system of ODEs, which are solved as an initial value problem, takes the form:

$$\frac{d\phi_i}{dt} = f(\phi, p_0), \quad i = 1, 2, 3, \dots, n \quad (2.1)$$

$$\phi(0) = \phi_0 \quad (2.2)$$

where ϕ is the thermo-chemical composition vector, which includes species mass fractions and temperature, the chemical source term $f(\phi, p_0)$ is a function of thermo-chemical scalars, ϕ , and is typically modeled using the law of mass action. The chemical kinetics presented by the above system of stiff ODEs is learned through the React-DeepONet framework of Kumar and Echehki (2024), where a reduced set of representative species is considered instead of the full thermochemical solution vector. This choice yields significant savings in computational time and memory, but accounting for a partial solution of the chemical system must also address mass and element conservation.

The dimensionality of the reduced composition space ($\phi_i, i = 1, 2, \dots, m \in \mathbb{R}^m$) is far smaller than the dimensionality of the original solution vector ($\phi_i, i = 1, 2, \dots, m, m+1, m+2, \dots, n \in \mathbb{R}^n$). In the Kumar and Echehki study (2024), the React-DeepONet that was used to capture the two-stage ignition process of n-dodecane at various initial conditions was implemented with seven representative species and temperature (i.e., $m = 8$), and the full mechanism required the integration of 348 species and temperature (i.e., $n = 349$).

It is important to note that for implementing the proposed approach in CFD calculations, the training data should be suitable for the targeted composition space in CFD. This data can be generated using 0D (Nguyen et al., 2021; Mehl and Aubagnac-Karkar, 2023) or 1D (e.g., 1D flamelets) reactors. Simpler 2D or 3D direct numerical simulations or a reduced scale or dimension represent alternatives to construct the targeted composition space (Owoyele and Echehki, 2017; Kumar et al., 2023).

The React-DeepONet architecture adopted by Kumar and Echehki (2024) is shown in Figure 1. It consists of two subnetworks, the Branch Net and the Trunk Net, which include as inputs the solution vector for the representative species and temperature at a time, t , and a prescribed time step, δt , on the




Figure 1. React-DeepONet: ANN for predicting the evolution of representative species and temperature.

Branch and Trunk nets, respectively. The outputs of the two subnetworks are combined in a concatenation layer, resulting in output for the solution vector for the representative species and temperature at a later time increment, $t + \delta t$. Based on its architecture, React-DeepONet maps the solution vector of a set of representative species and temperature at a time of the evolution of this solution across a time step δt , thus avoiding the costly direct chemistry integration and enabling the time step, δt , to be prescribed primarily by transport time scale requirements in reacting flow simulations. To train the React-DeepONet, the following loss function is evaluated, which compares the React-DeepONet predictions to the training data, which corresponds to the numerically integrated solution of the full chemical system:

$$\mathcal{L}_{\text{DATA}} = \frac{1}{m} \sum_{i=1}^m (\phi_i(t + \delta t) - \phi_i^*(t + \delta t))^2 \quad (2.3)$$

where δt varies from 0 to a large fixed time step Δt and $\phi_i^*(t + \delta t)$ is the corresponding ground-truth value. Therefore, React-DeepONet accommodates training for a range of values for δt within Δt . During prediction, any value of δt within this range can be used or adaptively changed during the simulation. The training prescribed above acts as a preconditioner for the model parameters, and additional refined training is implemented based on an autoregressive procedure (Kumar and Echekki, 2024). Within the process, the refined model is trained to predict the solution vector at multiple time steps, N_t . Accordingly, the loss function of the refined training model combines the loss functions corresponding to steps from 1 to N_t .

2.2. Mapping nonrepresentative species from the reduced set: Correlation Net

The objective of the proposed approach is twofold: (1) to constrain the React-DeepONet predictions to satisfy mass and element conservation and (2) to predict nonrepresentative species, which are not predicted using React-DeepONet or transported in CFD. These species, combined with the representative species, may account for the bulk of the mass in the mixture. Although the dimensionality of the full solution vector is very large, it resides on a low-dimensional manifold since most of the scalars are highly correlated (Pope, 2013; Coussement et al., 2016). The reduced set ($\phi_i, i = 1, 2, \dots, m$), identified based on principal component analysis (PCA) of reaction rate data as proposed by Alqahtani and Echekki (2021), is representative of the thermo-chemical manifold. It strongly correlates with the nonrepresentative species ($\phi_i, i = m + 1, m + 2, \dots, n$). However, other methods for selecting the representative species can be used including directed relations graphs (DRGs) methods (Lu and Law, 2005), the global pathway selection (GPS) algorithm (Gao et al., 2016), or the more recent manifold-informed state vector subset (Zdybał et al., 2023). Recently, we showed that the PCA-based technique and the other three techniques yield similar representative species for chemical systems (Gitushi and Echekki, 2024).

Leveraging their high correlation, nonrepresentative species can be mapped from the reduced set of representative species using an ANN denoted as Correlation Net as shown in Figure 2. The training of the ANNs is implemented using training data from simulations that capture the accessed composition space of the target simulations in CFD. Inputs and outputs to Correlation Net are the representative species and temperature and nonrepresentative species, respectively, as shown in Figure 2. The loss that is evaluated during training is expressed as

$$\mathcal{L}_{\text{C-DATA}} = \frac{1}{(m-n)} \sum_{i=m+1}^n (\phi_i - \phi_i^*)^2 \quad (2.4)$$

Again, the loss compares predictions by Correlation Net with values obtained from the training data and resulting from the direct integration of the full thermo-chemical scalars' solution vector. Since the full thermo-chemical space obeys certain physical laws, such as the sum of species mass fractions is one and




Figure 2. Correlation Net: ANN for recovering nonrepresentative species.

elements remain conserved during the reaction as the solution vector evolves, Correlation Net is trained with additional losses corresponding to these physical constraints.

$$\mathcal{L}_{\text{C-SUM-MF}} = \left| \sum_{i=2}^n \phi_i - 1.0 \right| \quad (2.5)$$

$$\mathcal{L}_{\text{C-ELEM}} = \sum_{K=1}^{N_{\text{Element}}} \left| \sum_{i=2}^n \frac{N_K^i W_K}{W_i} (\phi_i - \phi_i^*) \right|, \quad (2.6)$$

where N_K^i is the K th element contribution in i th species, W_K is the atomic weight of K th element, and W_i is the molecular weight of the i th species. Also, ϕ_1 is essentially temperature and is excluded from the conservation loss formulations. It should be noted that these two losses include the Correlation Net output and its input. The total loss function for Correlation Net takes this form:

$$\mathcal{L}_{\text{C}} = W_1 \mathcal{L}_{\text{C-DATA}} + W_2 \mathcal{L}_{\text{C-SUM-MF}} + W_3 \mathcal{L}_{\text{C-ELEM}}, \quad (2.7)$$

which combines losses associated with correlations and mass and element conservation. In this expression, W_1 , W_2 , and W_3 are the weights for each loss term. The choice of the weights is dictated by the magnitudes of the various loss terms within Eq. (2.7) as discussed below. Although mass and element conservation may be redundant for reactions, they represent two useful constraints to enforce during the Correlation Net training. The results discussed below will investigate whether implementing both constraints improves the training of Correlation Net.

It is important to note that alternative approaches have been developed in the literature to reinforce mass and element conservation within the framework of ANN-based predictions of chemistry. Readshaw et al. (2023) implemented mass conservation in different steps. In the first step, the species are predicted by training an ANN without mass conservation constraints. The subsequent steps are implemented outside the ANN predictions. In these steps, residual species representing the main contributions to the elements are determined based on the ANN predictions. Then, the mass fractions of these species are recalculated to conserve their corresponding elements across the integration time step. An additional step is implemented if the new value of the residual species is negative by setting its value to zero and enforcing element conservation using two rectifying species.

Mehl and Aubagnac-Karkar (2023) also adopted a correction scheme for element conservation on selected species. For H₂ oxidation, H and O elements are corrected by adding to or removing mass from H₂ and O₂ from the ANN predictions. They also implemented an element-based hybrid model to replace the ANN evaluation with direct integration predictions when an error proxy exceeded a threshold value. This error proxy is based on the relative errors achieved by the predictions of H and O elements' mass fractions during the reaction.

Rohrhofer et al. (2024) implemented mass conservation within the network architecture by adding a softmax or a Sigmoid activation function to the last layer of the network to predict species from an input represented by a mixture fraction, a progress variable, and a dissipation rate using flamelet libraries. In contrast to the above studies, the present study implements the physical constraints of mass and element conservation through the loss function. While no additional constraints are imposed beyond the training process, the present results indicate that imposing these constraints can still significantly improve the predictions of our solution mapping approach. More importantly, it provides, within the same loss formulation, the option to weigh in the species predictions versus their mass/element constraints differently.

2.3. Physics-constrained chemical kinetics identification

As stated earlier, the prediction of the representative species and temperature is implemented through a DeepONet. The nonrepresentative solution vector can be reconstructed through the pre-trained Correlation Net for any reduced solution vector. The mapping from Correlation Net extends the reduced set to the full solution vector as shown in Figure 3. This figure illustrates the prediction of the entire thermochemical state solution vector using React-DeepONet for temperature and representative species and augmented with a pretrained correlation net to recover nonrepresentative species while enforcing conservation principles through loss evaluations. Utilizing the mapping provided by pretrained Correlation Net, the training process of React-DeepONet can now be enriched with the physical constraints of species and elemental mass conservation, as formulated in Section 2, with the corresponding losses defined:

$$\mathcal{L}_{\text{SUM-MF}} = \left| \sum_{i=2}^n \phi_i(t + \delta t) - 1.0 \right| \quad (2.8)$$




Figure 3. Full ANN coupling of React-DeepONet and Correlation Net.

$$\mathcal{L}_{\text{ELEM}} = \sum_{K=1}^{N_{\text{Element}}} \left| \sum_{i=2}^n \frac{N_K^i W_K}{W_i} (\phi_i(t + \delta t) - \phi_i^*(t + \delta t)) \right| \quad (2.9)$$

in addition to the data loss that is defined as

$$\mathcal{L}_{\text{DATA}} = \frac{1}{m} \sum_{i=1}^m (\phi_i(t + \delta t) - \phi_i^*(t + \delta t))^2 \quad (2.10)$$

The total loss becomes

$$\mathcal{L} = W_1 \mathcal{L}_{\text{DATA}} + W_2 \mathcal{L}_{\text{SUM-MF}} + W_3 \mathcal{L}_{\text{ELEM}} \quad (2.11)$$

where W_1 , W_2 , and W_3 are the weights for each loss term. Again, element and mass conservation as in Eq. (2.7) may be redundant. Therefore, enforcing element conservation loss alone is sufficient to establish mass and element conservation, and either weight, W_1 or W_2 , can be set to zero. Although, preferably, W_2 should be kept and W_1 can be set to zero. We will again investigate whether enforcing this redundancy in mass and element conservation can improve the training in the Results section.

Finally, we note that the solution vector in the network shown in Figure 3 is subjected to both a transformation and a normalization. The main exception is where we implement the loss function evaluation, where we need to transform the species mass fractions to their actual values to impose mass and element conservation. The transformation and scaling procedure is implemented as follows:

- A fifth-order transformation is implemented on the species profiles and temperature to accommodate their strong variations. For example, for the k th species mass fraction, Y_k , this transformation is implemented as follows:

$$y_k = Y_k^{1/5} \quad (2.12)$$

- In a second step, we normalize the transformed variables, y_k , to yield values in the range of -1 to 1 among all the training data by using the minimum and maximum values of each transformed quantity.

3. Results

This section presents validation results for our proposed framework to predict species profiles and reduce enforced mass and element conservation. As stated above, we demonstrate the framework using our React-DeepONet approach (Kumar and Echehki, 2024). The data is generated using 0D homogeneous chemistry simulations at constant atmospheric pressure for CH_4 oxidation. We have chosen to validate our approach using 0D data and without advective or diffusive transport typical of 3D reacting flow simulations to unambiguously establish element and mass conservation with the proposed framework. Although Mehl and Aubagnac-Karkar (2023) relied on a model that involves 0D stochastic particles with mixing, the enforcement of mass conservation is implemented on the reaction source operator only. A similar strategy was adopted by Readshaw et al. (2023), who used unsteady premixed flames but enforced mass conservation on the reaction source operator. It is equally important to note that the proposed framework can be extended to a broad range of chemical states (including states that result from reaction and transport) by using multiple React-DeepONets with mass/element conservation for different regions of the composition space of interest using clustering techniques.

The solution profiles for training purposes are generated using Cantera (Goodwin et al., 2022) with the GRI-Mech 3.0 (Smith et al., 1999) mechanism with 53 species and 325 reversible reactions. The solution vector consisting of species mass fractions and temperature is evolved in time for a total of 100 initial conditions of methane–air mixtures with 10 initial temperatures equally spaced from $T_i = 1100$ K to 1150 K and 10 equivalence ratios equally spaced from 0.8 to 1.2. Of these 100 profiles, 90 are used for

training and 10 for testing. This data corresponds to outputs of the solution vector at increments of 10^{-6} s from 0 to 8 ms.

Instead of the full thermo-chemical vector (54 scalars), the reduced set of eight representative species, including H_2 , H_2O , CH_4 , CO_2 , CH_3OH , O_2 , CO , and NO , and temperature are considered for the chemical kinetics learning through the React-DeepONet model. The list of representative species is chosen based on the two-step PCA process of Alqahtani and Echekki (2021). The method is designed to select species whose reactions contribute the most to all species in the chemical system. The list also includes all the major species involved during methane oxidation. The remaining 45 species in the GRI-Mech 3.0 are modeled as nonrepresentative species within the proposed framework. Their mass fractions are modeled using Correlation Net based on the input of the temperature and eight representative species listed above.

Other methods can be used to determine representative scalars. They include the use of DRGs (Lu and Law, 2005), the GPS algorithm of Gao et al. (2016), or the manifold-informed representative species selection method by Zdybał et al. (2023).

As discussed above, Correlation Net is pretrained with the physical constraints on the full dataset, where the representative and nonrepresentative species are inputs and outputs to the model, respectively. Then, utilizing the mapping from the pretrained Correlation Net, the React-DeepONet model is trained with the physical constraints of mass and elemental conservation.

Both Correlation Net and React-DeepONet are modeled and trained with the JAX framework (Frostig et al., 2018) on an NVIDIA RTX A5500 GPU with the Adamax optimizer (Kingma and Ba, 2014). The size of trainable parameters of Correlation Net is around 2000, and for React-DeepONet, it is around 200,000. Correlation Net and React-DeepONet are trained for 300,000 and 80,000 iterations, taking around 3.3 minutes and 10.5 minutes, respectively, of GPU time. The detail of architectures and training hyperparameters for both React-DeepONet and Correlation Net are listed in Table 1. As the table shows, much of the network complexity resides in the React-DeepONet network, while the reconstruction of the nonrepresentative species is implemented with only two hidden layers. However, the total training time is comparable. Moreover, the same activation function, Tanh, and the same loss function, MSE, are used in both networks.

The obtained input–output pairs from training initial conditions are further broken down in an 80/20 ratio to extract the validation dataset. Figure 4 shows the decay of data loss, the sum of mass fraction loss, and elemental conservation loss during the training of Correlation Net and React-DeepONet for training as well as the validation dataset. Different observations can be made related to the trends of the different loss terms. First, the validation loss decay is almost identical between training and validation losses, indicating that the model hyperparameters are well optimized. Second, the mass and element conservation

Table 1. ML models and training hyperparameters

Parameters	React-DeepONet	Correlation Net
Architecture	Branch Net: [9,100,100,100,100,450] Trunk Net: [1,100,100,100,100,450] Param Net: [1,100,450]	[9,20,20,45]
Total parameters	200 K	2000
Loss function	MSE	MSE
Activation function	Tanh	Tanh
Training iterations	8×10^4	3×10^5
Training epochs	90	4160
Total training time (GPU minutes)	10.5	3.3
Mini-batch size	20,000	20,000
Learning rate	8×10^{-3}	8×10^{-3}

Note. The architecture of each network or subnetwork is indicated within square brackets, showing the number of neurons per layer.




Figure 4. Loss decay for training and validation dataset. (a) Correlation Net training loss decay with physical constraints of species and elemental mass conservation. (b) React-DeepONet training loss decay without physical constraints. (c) React-DeepONet training loss decay with physical constraints.

losses exhibit the same trends in their rate of convergence and magnitude in the Correlation Net and the combined network losses. Note that the weights for these losses are assigned the same values. Therefore, as suggested above, the element and mass conservation losses are redundant, and enforcing mass and element conservation simultaneously is not needed. Therefore, we can eliminate considerations for the mass conservation losses in both the pretrained Correlation Net and the combined networks. Finally, by adding conservation losses (see Figure 4), the errors associated with these losses are significantly reduced relative to the data loss contributions.

In Figure 4a, the decay in the loss corresponding to the species and elemental mass conservation in the training of Correlation Net ensures that the trained model is physically more compliant. Figure 4b,c shows the React-DeepONet training loss with and without constraints of conservation laws, respectively. It is evident that, during the kinetics learning from React-DeepONet, losses corresponding to the conservation laws further decay by more than one order of magnitude when corresponding conservation constraints are imposed during the training. Figure 4 suggests that a range of values for the weights in Eq. (2.7) can be adapted to highlight a particular loss contribution. With equal weights, the dominant loss terms may correspond primarily to those of the solution mapping or the sum of the mass fractions constraints. In this study, we selected the following weights for individual losses in the total loss formulation: W_1 , W_2 and W_3 are 0.6, 0.2, and 0.2, respectively, for the training of both Correlation Net and React-DeepONet. We find that a range of values of these weights, within a factor of 2 or 3, yields similar outcomes in mapping solution vectors across time steps while enforcing mass conservation constraints.

3.1. Validation of Correlation Net

In this section, we first validate the standalone Correlation Net before its integration with React-DeepONet. In this validation study, we construct the Correlation Net model using the Cantera simulations on test data, where nonrepresentative species (the output of Correlation Net) are evaluated in terms of the representative species (the input of Correlation Net). The comparisons are carried out using the Cantera predictions for the nonrepresentative species. Individual temporal profiles of species are compared for one run condition, corresponding to an initial methane-air mixture at 1135 K and an equivalence ratio of 1.1. This condition is part of the testing data, not the training data. Similar trends are observed for other initial temperatures and mixture conditions. We will also present comparisons of the different test cases by integrating errors across the entire time evolution of individual profiles.

Figure 5 compares the temporal profiles of three nonrepresentative and minor species, OH, C_2H_4 , and NH_3 , mass fractions based on the Cantera solutions (solid lines) and their recovery using




Figure 5. (Correlation Net) Reconstruction of three nonrepresentative species temporal profiles for $T_i = 1135$ K and an equivalence ratio of 1.1.

Correlation Net (symbols). Note that the solutions are normalized using minimum and maximum values to illustrate the results for the various species in one plot. The results show an initial methane–air mixture at 1135 K and an equivalence ratio of 1.1. They correspond to the case where mass and element conservation are integrated with the loss evaluations of Correlation Net. The comparison of the three nonrepresentative species shows an excellent agreement, indicating that Correlation Net provides an adequate prediction of nonrepresentative species. Similar results are obtained for other nonrepresentative species in the methane mechanism but are not shown here. More importantly, implementing mass and element conservation constraints significantly reduces the prediction errors of Correlation Net.

Figure 6 shows comparisons of the temporal profiles of the NH_3 mass fraction from the Cantera solution (blue lines) and the Correlation Net predictions with (green lines) and without (red lines) mass and element conservation enforced. The figure also illustrates the errors in the predictions of the mass and element conservation with and without the mass/element conservation constraints. For mass conservation, the reported value corresponds to the sum of the mass fractions of the representative and nonrepresentative species, recovered by Correlation Net. With mass conservation, this sum should have a value of 1. For element conservation, the results are shown for the elemental mass fraction, which can be expressed for the element i as follows:

$$y_i = \sum_{k=1}^N \mu_{ik} Y_k, \quad (3.1)$$

where y_i is the elemental mass fraction for element i , Y_k is the mass fraction of the k th species, and μ_{ik} is the mass fraction of element i in the k th species. This elemental mass fraction must remain constant for homogeneous chemistry since reactions do not alter elements. Moreover, the sum of all elemental mass fractions must also be 1.

The NH_3 mass fraction profiles for the Correlation Net with and without conservation constraints are indistinguishable. The difference between the two sets of predictions can only be seen when looking closely at the errors incurred for mass and element conservation. From the profiles, it is clear that imposing these errors is much lower when mass and element conservation are incorporated within the loss evaluations of Correlation Net.

As expected, the elemental mass fractions from the Cantera solution remain fixed for all elements considered. With Correlation Net, incorporating the mass and element conservation losses significantly reduces the deviation from the expected constant predictions. Yet, this deviation remains reasonably small whether conservation constraints are imposed or not.

So far, we have presented the above results based on one test case corresponding to a prescribed initial temperature and equivalence ratio. Figure 7 shows the mean absolute errors (MAEs) associated with mass and element conservation for the 10 test cases considered using Correlation Net with and without




Figure 6. (Correlation Net) Prediction of NH_3 profiles and validation of mass and element conservation with and without physical constraints ($T_i = 1135$ K and an equivalence ratio of 1.1).

incorporating conservation losses. The MAEs are based on averages across different states along the integration trajectory for each solution using the original denormalized quantities. For mass conservation, the MAE is based on the deviation of the sum of representative and nonrepresentative species relative to 1. For element conservation, the MAE is based on the deviation of the elemental mass fraction from the constant value computed in Cantera for each element.

The results clearly show that Correlation Net with and without the incorporation of losses can still result in errors that are less than 0.1% of the fixed values for the sum of mass fractions or elemental mass. As expected, incorporating mass and element conservation significantly reduces the error of recovering nonrepresentative species with Correlation Net. This reduction in error is approximately one order of magnitude for C, H, O, and N.

Figure 8 shows the MAE for species and temperature predictions for all test cases based on the training of Correlation Net and corresponding to the scenarios of implementing mass/element constraints or not implementing them. Recall that both scenarios are trained with an MSE loss formulation.




Figure 7. (Correlation Net) Time-averaged mean absolute errors for mass and element conservation with and without conservation constraints for all test cases.




Figure 8. (Correlation Net) Time-averaged MAE for species and temperature predictions from Correlation Net with and without constraints for all the test cases.

The reconstruction MAE tends to be slightly higher with the implementation of the constraints, given the additional contributions to the loss function. Nonetheless, the predictions tend to be uniform across all test cases.

Finally, it is important to note that the Correlation Net predictions of the species presented in Figure 5 are typical for all solution profiles within the range of the training data and for all species evaluated through Correlation Net. No discernable variations exist between OH, C₂H₄ and NH₃ predictions and those of other species in the mechanism. Therefore, the time-averaged MAE plots shown in Figure 8 represent the performance of the individual species reconstructed using Correlation Net.

3.2. *A posteriori comparisons with Correlation Net and React-DeepONet*

Next, we present the results of the a posteriori predictions of Correlation Net with and without mass and element conservation with its integration of React-DeepONet. In contrast with the previous discussion, the a posteriori calculation of the species involves the evaluation of the evolution of the representative species in time and is implemented using React-DeepONet. The nonrepresentative species are again recovered using Correlation Net with an input corresponding to the predicted representative species.

Figures 9–11 illustrate the same results presented in Figures 5–7, respectively. As shown in Figure 9, the a posteriori results for combining Correlation Net and React-DeepONet can accurately reproduce profiles of nonrepresentative species. Similar to earlier observations, Figure 10 also shows that adding mass and element conservation in Correlation Net achieves its desired objectives. This is demonstrated in one case in Figure 10 in addition to 10 cases, as shown in Figure 11. We observe similar trends for the MAE error reduction for mass and element conservation as obtained using Correlation Net alone.

Figure 12 shows the time-averaged MAE for species and temperature predictions from the entire network combining React-DeepONet and Correlation Net for all test cases considered for the two scenarios, where mass/element conservation is implemented or not. Again, both scenarios are trained with an MSE loss formulation. In contrast with the predictions based on Correlation Net alone, as shown in Figure 8, the MAE values can be higher or lower between the two scenarios. Yet, their magnitudes are again comparable and do not vary significantly among the different test cases.

4. Conclusions

In this study, we propose and implement a physics-constrained model for mass and element conservation within a framework of chemistry integration and an approach to recover nonrepresentative species while advancing the solution of a subset of the chemical system. The implementation involves the coupling of




Figure 9. (Correlation Net + React-DeepONet) Reconstruction of three nonrepresentative species temporal profiles for $T_i = 1135$ K and an equivalence ratio of 1.1.




Figure 10. (Correlation Net + React-DeepONet) Prediction of NH_3 profiles and validation of mass and element conservation with and without physical constraints ($T_i = 1135$ K and an equivalence ratio of 1.1).

the reduced-state framework with a Correlation Net to recover nonrepresentative species and constraints on mass and element conservation implemented on the overall network. The coupled framework significantly improves mass and element conservation and can predict (representative and nonrepresentative) species as a potentially reduced computational cost.

The principal advantage of the proposed approach is the potential implementation of the chemistry integration framework with element and mass conservation in reacting flow simulations. With a choice of a subset of the chemical species and mapping solutions across time increments, computational and memory savings can be achieved. The advantage of this framework for complex combustion fuels can be critical as the solution vector can be reduced from hundreds or thousands of transported species to a range from 10 to 20 representative species. The proposed scheme is also able to recover nonrepresentative species. With the use of Correlation Net, the nonrepresentative species can be used for postprocessing purposes in addition to the accurate evaluation of thermodynamic and transport properties during a simulation.




Figure 11. (Correlation Net + React-DeepONet) Time-averaged mean absolute errors for mass and element conservation with and without conservation constraints for all test cases.




Figure 12. (Correlation Net + React-DeepONet) Time-averaged MAE for species and temperature predictions from React-DeepONet with and without constraints for all the test cases.

Data availability statement. For any queries related to the data used and analyzed in this study, please contact the corresponding author (Tarek Echekki, techekk@ncsu.edu). The code is available at: <https://github.com/TEchekki/Physics-Constrained-Machine-Learning-for-Reduced-Composition-Space-Chemical-Kinetics>.

Author contribution. Conceptualization: A.K., T.E.; Methodology: A.K., T.E.; Data curation: A.K.; Writing original draft: A.K.; Draft editing: A.K., T.E. All authors approved the final submitted draft.

Funding statement. This research was supported by grants from Aramco-Americas through a contract number ASC Agreement No. CW47993.

Competing interests. The authors declare none.

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

References

- Alqahtani S and Echekki T (2021) A data-based hybrid model for complex fuel chemistry acceleration at high temperatures. *Combustion and Flame* 223, 142–152.
- Blasco J, C D, Fueyo N and Chen J-Y (2000) A self-organizing-map approach to chemistry representation in combustion applications. *Combustion Theory and Modelling* 4, 61–76.
- Blasco J, Fueyo N, Dopazo C and Ballester J (1998) Modelling the temporal evolution of a reduced combustion chemical system with an artificial neural network. *Combustion and Flame* 113, 38–52.
- Coussement A, Isaac BJ, Gicquel O and Parente A (2016) Assessment of different chemistry reduction methods based on principal component analysis: Comparison of the MG-PCA and score-PCA approaches. *Combustion and Flame* 168, 83–97.
- Dikeman HE, Zhang H and Yang S (2022) Stiffness-reduced neural ode models for data-driven reduced-order modeling of combustion chemical kinetics. In: *AIAA Scitech 2022 Forum*. Paper Number AIAA 2022-0226. <https://doi.org/10.2514/6.2022-0226>
- Frostig R, Johnson M and Leary C (2018) *Compiling Machine Learning Programs via High-Level Tracing*. Available at <https://mlsys.org/Conferences/doc/2018/146.pdf>
- Gao X, Yang S and Sun W (2016) A global pathway selection algorithm for the reduction of detailed chemical kinetic mechanisms. *Combustion and Flame* 167, 238–247.
- Gitushi KM and Echekki T (2024) Comparisons of different representative species selection schemes for reduced-order modeling and chemistry acceleration of complex hydrocarbon fuels. *Energies* 17, 2604.
- Goodwin DG, Moffat HK, Schoegl I, Speth RL and Weber BW (2022) Cantera: An object-oriented software toolkit for chemical kinetics, thermodynamics, and transport processes [Version 2.6.0].
- Goswami S, Jagtap AD, Babae H, Susi BT and Karniadakis GE (2024) Learning stiff chemical kinetics using extended deep neural operators. *Computer Methods in Applied Mechanics and Engineering* 419, 116674.
- Haghshenas M, Mitra P, Santo ND and Schmidt DP (2021) Acceleration of chemical kinetics computation with the learned intelligent tabulation (LIT) method. *Energies* 14, 7851.
- 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* 125, 1082–1092.
- Kingma DP and Ba J (2014) Adam: A method for stochastic optimization. *Proceedings of the 3rd International Conference on Learning Representations (ICLR)*, Banff, 14–16 April 2014. [arXiv:1412.6980](https://arxiv.org/abs/1412.6980).
- Kumar A and Echekki T (2024) Combustion chemistry acceleration with deeponets. *Fuel* 365, 131212.
- Kumar A, Rieth M, Owoyele O, Chen JH and Echekki T (2023) Acceleration of turbulent combustion dns via principal component transport. *Combustion and Flame* 255, 112903.
- Lu L, Jin P, Pang G, Zhang Z and Karniadakis GE (2021) Learning nonlinear operators via deepnet based on the universal approximation theorem of operators. *Nature Machine Intelligence* 3, 218–229.
- Lu T and Law C (2005) A directed relation graph method for mechanism reduction. *Proceedings of the Combustion Institute* 30, 1333–1341.
- Mehl C and Aubagnac-Karkar D (2023) On-the-fly accuracy evaluation of artificial neural networks and hybrid method to improve the robustness of neural network accelerated chemistry solving. *Physics of Fluids* 35(6), 067115.
- Nguyen H-T, Domingo P, Vervisch L and Nguyen P-D (2021) Machine learning for integrating combustion chemistry in numerical simulations. *Energy and AI* 5, 100082.
- Owoyele O and Echekki T (2017) Toward computationally efficient combustion dns with complex fuels via principal component transport. *Combustion Theory and Modeling* 21, 770–798.
- Pope SB (2013) Small scales, many species and the manifold challenges of turbulent combustion. *Proceedings of the Combustion Institute* 34, 1–31.
- Raissi M, Perdikaris P and Karniadakis G (2019) Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics* 378, 686–707.

- Ranade R, Alqahtani S, Farooq A and Echekki T** (2019) An ANN based hybrid chemistry framework for complex fuels. *Fuel* 241, 625–636.
- Readshaw T, Jones WP and Rigopoulos S** (2023) On the incorporation of conservation laws in machine learning tabulation of kinetics for reacting flow simulation. *Physics of Fluids* 35(4), 047103.
- Rohrhofer FM, Posch S, Gößnitzer C, García-Oliver JM and Geiger BC** (2024) Utilizing neural networks to supplant chemical kinetics tabulation through mass conservation and weighting of species depletion. *Energy and AI* 16, 100341.
- Smith GP, Golden DM, Frenklach M, Moriarty NW, Eiteneer B, Goldenberg M, Bowman CT, Hanson RK, Song S, William C Gardiner J, Lissianski VV and Qin Z** (1999) GRI-Mech. http://www.me.berkeley.edu/gri_mech/.
- Venturi S and Casey T** (2023) SVD perspectives for augmenting DeepONet flexibility and interpretability. *Computer Methods in Applied Mechanics and Engineering* 403, 115718.
- Zdybal K, Sutherland JC and Parente A** (2023) Manifold-informed state vector subset for reduced-order modeling. *Proceedings of the Combustion Institute* 39, 5145–5154.