Session 6

Combustion, Flames, Sprays

S6-A2 Cédric Mehl — Acceleration of chemical kinetics integration in combustion CFD using artificial neural networks

S6-A4 Min Xu — High-Spatiotemporal-Resolution Combustion Diagnostics using AI-based Reconstruction Techniques

S6-A5 Nicolas Eckel — ANN-based Chemistry Surrogate Modeling for Premixed Hydrogen Flames

S6-A6 Cyril Freyling — Symbolic Regression-Based Prediction of Mass Fraction Burned in Hydrogen-Fueled Internal Combustion Engines

S6-P1 Baptiste Romain — Chemistry reduction and acceleration based on neural network surrogate

Acceleration of chemical kinetics integration in combustion CFD using artificial neural networks

[1]C. MEHL*

[1]*Institut Carnot IFPEN Transports Energie, IFPEN, 1 et 4 av. du Bois Préau, Rueil-Malmaison, 92852, France.*

Abstract

The numerical integration of chemistry is often the bottleneck of reactive fluids simulations, in particular in combustion CFD. A solution that has attracted a lot of attention in the past years is to use artificial neural networks (ANN) to replace the costly but accurate implicit solvers. The main issue is then to generate databases relevant to the targeted simulation configurations and ensure the robustness and accuracy of the ANN based integration. A strategy has recently been proposed at IFPEN featuring the use of 0D stochastic reactors to define the training database and a hybrid ANN model for chemistry integration. The main idea is to switch to the implicit solver in regions where the accuracy of the ANN is deemed insufficient, based on a metric built using mass conservation. In this work, we present the latest advances of the hybrid ANN strategy developed at IFPEN. It features more specifically (1) the extension of the model to multi-time step prediction; (2) the validation of the strategy on an academic swirl burner fed with NH3/H2.

Keywords: combustion, neural networks, chemistry integration

High-Spatiotemporal-Resolution Combustion Diagnostics using AI-based Reconstruction Techniques

Prof. Min Xu, Prof. Xuesong Li, Shanghai Jiao Tong University, China

Abstract:

Quantitative measurement of three-dimensional (3D) combustion flame fields is pivotal for investigating the spatial topological structures and transient evolutionary mechanisms of turbulent combustion. While tomography serves as a primary tool for 3D optical diagnostics, conventional methods rely heavily on spatial discretization grid resolution. Consequently, reconstruction quality is constrained by the inherent trade-off between discretization accuracy and computational complexity. These methods struggle to simultaneously achieve sub-millimeter spatial resolution and kilohertz-level temporal resolution, thereby severely limiting the capability to resolve fine-scale flame structures.

To address these challenges, this study proposes a novel method for high-spatiotemporal-resolution 3D flame reconstruction based on Neural Implicit Representation. Unlike traditional discrete voxel-based Algebraic Reconstruction Techniques (ART), the proposed method utilizes Multi-Layer Perceptrons (MLPs) to model the 3D combustion field as a continuous coordinate mapping function. This approach marks a paradigm shift from discrete grid descriptions to continuous full-field representations. Spatially, the method employs coordinate networks to generate a voxel-free implicit 3D representation of transient flames. By integrating adaptive spatial sampling strategies for regions of interest (ROI), it overcomes the physical resolution limits of imaging hardware, achieving sub-millimeter spatial super-resolution. Temporally, we construct an implicit 4D spatiotemporal representation model for dynamic 3D flames. By encoding time as a continuous variable, a single network model can perform high-precision fitting and interpolation across thousands of frames in dynamic flame sequences, effectively capturing the dynamic stretching and wrinkling characteristics of flame surfaces.

Numerical simulations and experimental validation demonstrate that the method maintains exceptional reconstruction fidelity even under sparse-view conditions. It significantly outperforms traditional tomographic techniques in terms of spatiotemporal resolution, noise robustness, and measurement dynamic range. This study provides a novel, intelligent solution to the persistent challenges of low resolution ("inability to see

clearly") and low frame rates ("inability to keep up") in complex fluid experiments. It holds significant application value for fundamental combustion research and diagnostics in practical engine combustion development.

Keywords:

Combustion diagnostics; 3D reconstruction; Neural Radiance Fields (NeRF); Implicit representation; High spatiotemporal resolution

ANN-based Chemistry Surrogate Modeling for Premixed Hydrogen Flames

[1]N. Eckel*; [1]T.J.P. Karpowski; [2]D. Mira; [3]F. Ferraro; [1]C. Hasse; [1]A. Scholtissek

[1]*Institute for Simulation of Reactive Thermo-Fluid Systems, TU Darmstadt, Otto-Berndt-Straße 2, 64287, Darmstadt, Germany*

[2]*Barcelona Supercomputing Center, Plaça d'Eusebi Güell, 1–3, Barcelona, 08034, Catalonia, Spain*

[3]*Institute of Jet Propulsion and Turbomachinery, TU Braunschweig, Hermann-Blenk-Str. 37, 38108, Braunschweig, Germany*

Abstract

In recent years, artificial neural network (ANN) based chemistry surrogate modeling has emerged as a promising paradigm for reducing computational cost in CFD simulations of reacting flows. In this approach, an ANN is first trained on a suitable thermochemical flame dataset, learning the intrinsic low-dimensional manifold. It is then coupled to a CFD solver, where it replaces the conventional numerical chemistry integration and provides thermochemical quantities. While significant progress has been made and several successful applications of this method have been demonstrated (by means of *a posteriori analyses*), they often rely on problem-specific network architectures and training approaches. A universally applicable methodology has yet to be established.

This work aims to improve the understanding of the challenges associated with training accurate surrogate models on multiscale combustion data. To this end, we identify relevant challenges, analyze them, and explain their connection to the combustion physics. We then introduce a methodology that systematically addresses these challenges. The method is used to developed ANN-based surrogate models for 1D freely-propagating H_2 flames. Both a priori and a posteriori results are compared with detailed chemistry (DC) reference solutions. Additionally, we train a model on lean thermodiffusively unstable H_2 flame data for which a posteriori analyses showcases generalization of the method to more complex combustion physics.

Our results indicate that, due to the multiscale nature of combustion data, both data normalization and data sampling are imperative for the training of accurate ANN-based surrogate models. Flame-independent methodologies for both normalization and sampling are presented. The resulting surrogate models show good agreement for global flame characteristics (flame speeds and fuel consumption rates) as well as local flame characteristics (radical species profiles) for all analyzed flames.

Keywords: combustion, machine learning, artificial neural networks, thermodiffusive instabilities, chemistry surrogate modeling

Symbolic Regression-Based Prediction of Mass Fraction Burned in Hydrogen-Fueled Internal Combustion Engines

[1]Freyling Cyril *; [1]Coussement Axel

[1]*Aero-Thermo-Mechanics Department, Université Libre de Bruxelles, Belgium*

Abstract

This study proposes a predictive combustion modeling framework for hydrogen-fueled internal combustion engines based on symbolic regression. A two-zone zero-dimensional thermodynamic model is employed, in which the combustion process is described through the evolution of the mass fraction burned (MFB). The MFB is modeled using the Wiebe function, an analytical S-shaped formulation that represents the cumulative fraction of fuel burned as a function of crank angle. The Wiebe function is characterized by shape parameters and combustion duration, which are traditionally determined through empirical fitting to experimental pressure data. In the present work, symbolic regression is used to identify explicit analytical relationships linking these parameters to key engine operating variables, including equivalence ratio, intake pressure, compression ratio, engine speed, and spark advance. This approach enables direct prediction of both the mass fraction burned and the corresponding mass burning rate without case-specific calibration, while preserving physical interpretability and predictive capability across varying operating conditions.

Keywords: Hydrogen engine, Machine learning, Symbolic regression, Mass fraction burned (MFB), Two zone model

Chemistry reduction based on neural network surrogate: comparison to DRGEP and SARGEP methods

Baptiste Romain[a,b,*], Adrien Langenais[a], Joël Dupays[a], Laurent Catoire[b]

[a] *DMPE, ONERA, Université Paris Saclay, Palaiseau 91120, France*
[b] *ENSTA, Institut Polytechnique de Paris, Palaiseau 91120, France*

Abstract

A neural network (NN) has been utilized to reproduce chemical simulations obtained with the methane/air GRI-3.0 mechanism in a CFD solver. The goal is to investigate our ability to capture soot precursors dynamics. The training dataset contained samples drawn from homogeneous reactor (0D) simulations that were used to reproduce auto-ignition processes in a homogeneous domain (pseudo 0D). After the network including all GRI-3.0 species was validated, a reduction method was introduced. It consisted in processing the dataset so that the network learns the evolution of a selected subset of species, while representing all discarded one by a single fictitious species that is transported by the CFD solver. The reduced network was validated on 0D reactors, where it accurately reproduced the dynamics of all species. Two different methods were then tested to replace the traditional thermodynamics solver of a one-dimensional CFD code and make it compatible with the introduction of a fictitious species. In this configuration, the network captured the evolution of the tracked soot precursor (C_2H_2) with higher accuracy than a reduced chemical mechanism produced by the conventional DRGEP or SARGEP reduction methods. Moreover, its precision was comparable to that of the full detailed mechanism, while offering computation time saving. These results demonstrate that this innovative reduction method can preserve fidelity in predicting minor species mass fraction and remain efficient for CFD integration.

Keywords: Chemistry Reduction; Machine Learning; Soot Precursors; Neural Network; CH_4/O_2 Combustion

*Corresponding author.

1. Introduction

Accurate prediction of soot precursors in hydrocarbon based liquid rocket engines is a key challenge for assessing ecological impact of space launchers [1]. High-fidelity chemical kinetics models provide detailed description of the formation and consumption of minor species, but the high number of simulated reactants makes them prohibitive for use in computational fluid dynamics (CFD) codes. Conventional reduction methods like Directed Relation Graph with Error Propagation (DRGEP) [2] attempt to address this limitation by removing the least important reactions and species while preserving accuracy on chosen quantities. However, those reduced mechanisms often fail to capture the behavior of soot precursors [3], that are essential for predicting pollutant formation as they are themselves minor species of low impact on the global reaction.

Recently, neural networks (NN) have emerged in the field of solving ordinary differential equations (ODEs) and thus in the field of combustion simulations [4–6]. By learning the time evolution of the thermochemical state, NN can replace traditional solvers, bypassing the need for stiff integration and drastically reducing computational costs especially when GPU acceleration is available [7].

Most research efforts are aiming at efficiently generating databases for NN to learn the broadest possible chemical space at a reasonable cost.

Another key challenge is to enforce physical constraints on the network predictions. It can be trained specifically to consider important constraints like conservation of chemical elements [8]. This is often referred to as a soft constraint because the output can still marginally violate physical laws. Another method, referred to as a hard constraint, consists in projecting the network output or rebalancing atomic species to strictly enforce physical laws [8, 9].

In this paper, the focus is on proposing a reduction method relying on a NN-based chemistry solver that is compatible with a CFD code. As building a general database is not the main focus here, we use a simple 0D database to validate the coupling of the network surrogate to an in-house one-dimensional CFD solver.

2. Neural network-based reduction of the full mechanism and implementation into a CFD solver

2.1. Learning database

A NN surrogate was developed to emulate detailed chemical kinetic within a CFD solver. Accurate predictions require the NN to learn chemical dynamics from a dataset representative of the chemical states that will be encountered in the simulations to be done. In this study, the database was generated using zero-dimensional auto-ignition simulations following the approach of Chen [10]. Although limited to homogeneous domains, this configuration is sufficient to develop and validate the proposed reduction methodology.

Each entry of the database corresponds to an input-output pair composed of the thermochemical state vector $\mathbf{X}_t = (T, P, Y_1, \ldots, Y_n)$ and its evolution after a fixed time step $\mathbf{X}_{t+\Delta t}$. By learning this mapping, the NN replaces traditional implicit chemistry integrators whose numerical cost scales with the size of $\mathbf{X}_t$ to the power of two or three depending on the algorithm.

A total of 1000 auto-ignition simulations of a methane/air mixture were performed using the GRI-3.0 mechanism. Initial conditions were randomly sampled from a Halton sequence [11] within the following hypercube: $T_{ini} \in [1600, 1800]\,\mathrm{K}$, $\phi_{ini} \in [0.7, 1.5]$, and $P_{ini} = 101325\,\mathrm{Pa}$ such that the mole number of O_2, N_2 and CH_4 is expressed as: $n_{CH_4} = 1$, $n_{O_2} = 2/\phi_{ini}$, $n_{N_2} = 7.52/\phi_{ini}$. The timestep Δt used for generating the database was set to $1\,\mu s$.

A uniform sampling of the simulated trajectories would have overrepresented burnt and unburnt states. Furthermore, it would not sample finely enough zones of the chemical manifold exhibiting high curvature. We thus used to our advantage the adaptive time stepping capability of the CVODE integrator to over-sample regions of rapid chemical change, resulting in a more balanced dataset. The final database contained approximately 1.68 million thermochemical states

2.2. Model description and reduction methodology

The NN architecture is based on the ResNet architecture [12] as described in Figure 1. Each layer of the residual blocks contains 350 neurons while the backbone and output layers contain 54 neurons. The network is optimized using the Adam optimizer [13], 2000 warm-up steps and a learning rate of 5e-3 decaying exponentially with a decay rate of 0.92 and 1500 decay steps. The input vector $\mathbf{X}_t$ was scaled differently depending on the physical meaning of the data. Temperature and pressure where standardized to zero mean and unit variance while the logarithm of mass fractions was computed before the zero mean, unit variance standardization. The database is split such that 25 % of the samples are used for validation.

To improve the final accuracy of the surrogate, a mixture of experts method has been used. The dataset was partitioned into five clusters using the K-means algorithm [14], and five different NN were trained, one for each cluster, following what is described in Chen *et al.* [15]. The reduction method consisted in aggregating minor species mass fractions into a single fictitious species Y_{fict}. Each reduced state vector $(T, P, \mathbf{Y}_{red}) = (T, P, Y_{a_1}, \ldots, Y_{a_k}, Y_{fict})$ was associated with mixture-averaged properties—specific heat c_p, enthalpy h and molecular weight $\overline{W}$—computed from the detailed composition during database generation. Two networks where thus trained and used into the CFD solver instead of the classical thermochemical solver,

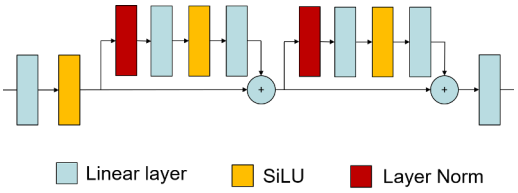

Fig. 1: Schematic view of the network architecture

namely: a chemistry NN (ChNN) to predict $(T, P, \mathbf{Y}_{red})_{t+\Delta t}$ and a thermodynamics NN (ThNN) to provide $c_p(T, \mathbf{Y}_{red})$, $h(T, \mathbf{Y}_{red})$, $\overline{W}(\mathbf{Y}_{red})$.

2.3. Coupling with the CFD solver

The code used for validation is a one-dimensional reactive Euler solver based on Strang splitting [16], where transport and chemistry are solved sequentially:

$$U^{n+1} = \mathcal{T}(\frac{\Delta t}{2})\mathcal{C}(\Delta t)\mathcal{T}(\frac{\Delta t}{2})$$

where, $\mathcal{T}$ and $\mathcal{C}$ are the transport and chemistry operators respectively.

In the baseline solver, chemical integration $\mathcal{C}$ is handled by a stiff ODE integrator based on an implementation of the backward differentiation formula [17]. This chemical integration operator $\mathcal{C}$ can also be replaced by the ChNN, which directly predicts the updated thermochemical state after a timestep Δt. All other parts of the numerical structure remains unchanged.

At each sub-step, conservative variables (ρ, ρu, ρE, ρY_i) are updated and primitive variables (T, P) are recomputed *via* thermodynamics properties of the mixture to ensure energy consistency.

Thermodynamics properties in the classical solver are evaluated from NASA-7 polynomials. However, since the reduced model introduces Y_{fict} with varying internal composition, using fixed NASA-7 coefficients will introduce an error. The ThNN thus replaces this evaluation, providing c_p, h and $\overline{W}$ as functions of $(T, \mathbf{Y}_{red})$.

Alternatively, an averaged NASA-7 entry for the pseudo-species has been introduced in the standard thermodynamics file, allowing direct use of the conventional solver to assess the error introduced by this approximation over using ThNN.

All the following results are obtained using this in-house CFD solver for simulating autoignitions.

3. Results

The performance of the proposed NN-based chemistry reduction was assessed against a baseline simulation using the detailed GRI-3.0 mechanism integrated by a conventional stiff ODE solver. For comparison, two established reduction techniques were also considered: Directed Relation Graph with Error Propagation (DRGEP) and Sensitivity Analysis-based Reduction assisted by Graph search with Error Propagation (SARGEP). Each approach was used to generate two mechanisms with the Brookesia software [18]: a soft reduction (tight accuracy constraints resulting in many kept species) and a heavy reduction (relaxed constraints, minimal species count). They were trying to reduce error generated on major species (CH_4, O_2, H_2O, CO_2) and on acetylene (C_2H_2). The resulting models are summarized in Table 1.

The ChNN was trained on a reduced set comprising 15 physical species (H_2, H, O, O_2, OH, H_2O, HO_2, CH_2, $CH_2(S)$, CH_3, CH_4, CO, CO_2, HCO, CH_2O, N_2) and one pseudo-species of mass fraction Y_{fict}. Its performance was evaluated on 400 unseen autoginition conditions.

To ensure consistent comparison, all reduced mechanisms were optimized at the reference conditions of $\phi = 1.0$ and $T = 1700\,K$, corresponding to the center of the NN training hypercube. The resulting ignition delays and species evolutions are listed in Table 1. From an accuracy standpoint, all models reproduce the autoignition delay well, except those heavily reduced by conventional methods, which lack intermediate species to resolve the oxidation sequence. We can also observe that using the ChNN with either NASA-7 thermodynamics or the ThNN gives the same result, while having a major computation efficiency difference, making of the ChNN with NASA-7 thermodynamics the preferred model.

The evolution of C_2H_2 mass fraction (Figure 2) highlights tree regimes:

- Soft conventional reductions maintain cor-

Table 1: Comparison of different thermochemical models in term of computational performance and precision.

Chemistry model	Thermodynamics model	Number of species	Chemistry time per iteration	Relative C_2H_2 peak error	Relative auto-ignition time error
GRI-3.0	NASA-7	53	0.352 s	—	—
NN with all species	NASA-7	53	0.023 s	4.17 %	0.92 %
ChNN with 16 species	NASA-7	16	0.016 s	3.31 %	0.92 %
DRGEP 24	NASA-7	24	0.150 s	17.7 %	298 %
DRGEP 35	NASA-7	35	0.160 s	67.7 %	1.37 %
SARGEP 28	NASA-7	28	0.230 s	12.7 %	0.46 %
SARGEP 15	NASA-7	15	0.141 s	100 %	56.9 %
ChNN with 16 species	ThNN	16	0.364 s	3.31 %	0.92 %

rect ignition timing but deviate moderately for acetylene prediction (10-20 %) while containing too many species for CFD simulations;

- NN-based models capture both the peak amplitude and timing within less than 5 % error;

- Heavy conventional reductions diverge significantly on every metrics (more than 50 %), with SARGEP 15 even failing to form acetylene.

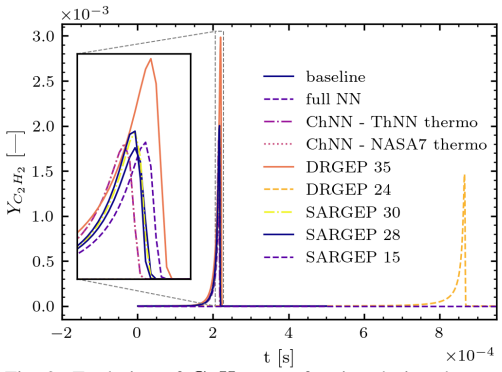

Fig. 2: Evolution of C_2H_2 mass fraction during the auto-ignition for every model.

Figure 3 to Figure 6 show two-dimensional maps of the mean absolute relative error of temperature and of the time integrated acetylene concentration, compared to detailed chemistry, as function of the initial temperature T_{ini} and equivalence ratio ϕ_{ini}. These maps directly illustrate each model's ability to reproduce the detailed chemical dynamics over a wide range of initial conditions.

For mean absolute error of temperature, both soft DRGEP and SARGEP maintain good accuracy near their design point but slightly degrade as mixture become fuel-rich. ChNN models show higher mean error but without clear spatial patterns else than increasing errors on the domain boundaries where training data is sparse. This behavior underlines a key limitation of data-driven models: while NNs generalize efficiently within their training manifold, extrapolation remains uncertain, making database coverage the true bottleneck of the method.

The integrated C_2H_2 concentration is directly relevant for soot prediction as soot models source terms are often proportional to precursors concentration [19]. Thus, Figure 6 provides more discriminating insight on the applicability of the method for pollutant predictions. DRGEP 24 model shows an error varying between 15 % and 45 % for temperatures above 1700 K. However, the localized nature of the reduction only preserves dominant pathways near the design point, while secondary routes that become major in off design conditions are neglected, leading to no soot formation for low temperature regions. SARGEP 28 performance, on the other hand, is comparable to the reduced neural network approach but retains twice as many species and incurs a significantly higher computational cost. Finally, hard reductions exhibit errors often above 50 % or fail to form acetylene. These results confirm the limited ability of graph-based reductions to reproduce soot precursor kinetics across broad thermochemical ranges while having a low species count.

In contrast NN-based reduced model with NASA-7 thermodynamics provides low error on the entire domain, typically below 6 % and often only a few percent. Despite being one of the models with the fewest species, it remains

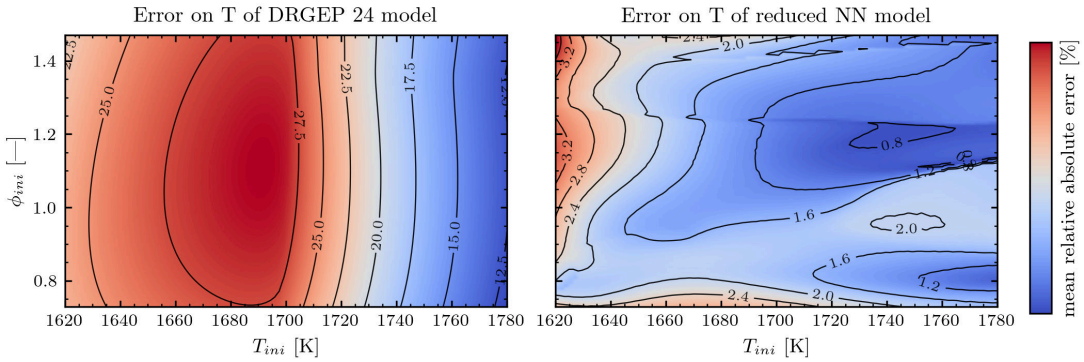

Fig. 3: Average relative absolute error compared to GRI-3.0 of temperature curves as a function of initial temperature and mixture ratio for DRGEP 24 model (left) and reduced NN model with NASA-7 thermodynamics model (right).

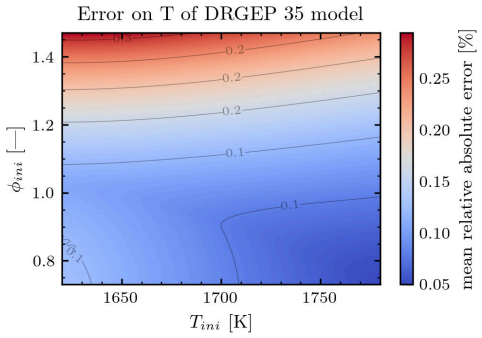

Fig. 4: Average relative absolute error compared to GRI-3.0 of temperature curves as a function of initial temperature and mixture ratio for DRGEP 35 model.

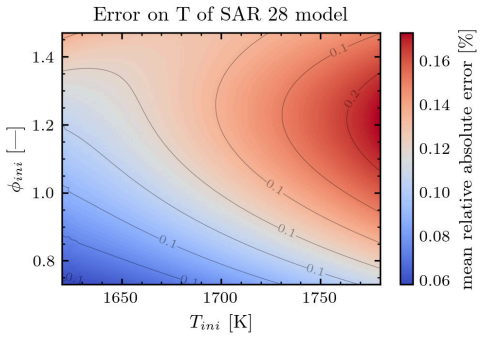

Fig. 5: Average relative absolute error compared to GRI-3.0 of temperature curves as a function of initial temperature and mixture ratio for SAR 28 model.

the most accurate and by far the most computationally efficient reduced model overall. This demonstrates the NN surrogate ability to capture nonlinear dynamics of soot precursor formation with a reduced species set and over a wide range of operating conditions, a key requirement for its later use in CFD soot modeling.

Computation times of traditional reduced models scales non linearly and rises quickly with the number of reactive species. However, replacing the stiff ODE integrator by a NN model reduces the computation time by a factor of 15.3, and an additional 43.7 % computation time is gained by using the reduced ChNN model.

Using the ThNN instead of NASA-7 polynomials, however, does not improve accuracy and drastically increases computation time, exceeding that of detailed integration, making it impractical in its current form. All benchmarks were performed on CPUs for consistent comparison. Nonetheless, inference tests on NVIDIA RTX A2000 GPU indicate an additional improvement in computation time by a factor of 2.5 to 3.3, underlying the strong potential of hybrid CPU/GPU configurations for accelerated combustion simulations.

4. Conclusion

This work demonstrates that NN-based chemistry reduction can reproduce the detailed chemical dynamics of soot precursors over a broad thermochemical range, while drastically reducing the computational cost. The NN surrogates outperform conventional reduction methods in its ability to use as few species as possible to re-

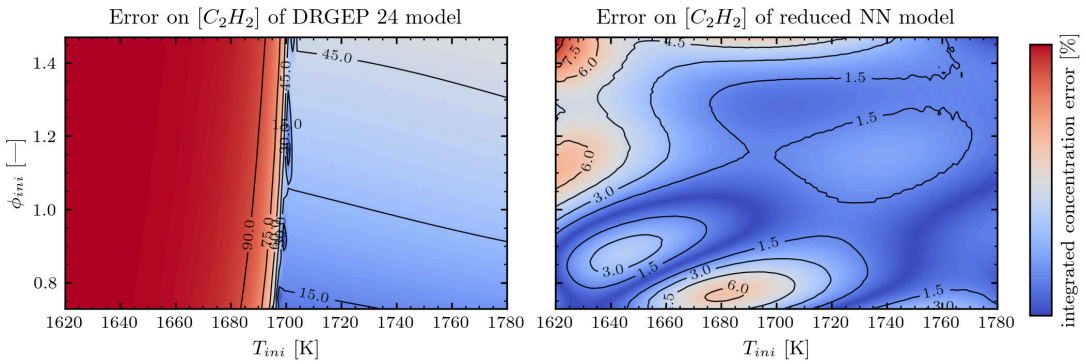

Fig. 6: Error on integrated C_2H_2 concentration compared to GRI-3.0 as a function of initial temperature and mixture ratio for DRGEP 24 model (left) and reduced NN model with NASA-7 thermodynamics model (right).

main precise on precursor kinetics such as acetylene, relevant for soot modeling.

Future developments should focus on optimizing the associated thermodynamics model. The averaged NASA-7 model gives great result if Y_{fict} remains small but it could be interesting to improve NN representation of thermodynamics for cases where Y_{fict} becomes larger.

In addition, the selection of species included in the reduced mechanism remains a key aspect. Introducing multiple pseudo-species, each designed to reflect species involved in a specific reaction pathway could enhance precision and allows for an even more reduced state.

Several directions are identified for future developments. First, the proposed reduction methodology could be applied to more complex chemical mechanisms, such as detailed hydrocarbon schemes with more soot precursors. Such mechanisms, which involve a significantly larger number of species and reactions, represent a much harder challenge for conventional graph-based reduction methods, as preserving both major species dynamics and minor pollutant precursors simultaneously becomes increasingly difficult. In this context, the NN-based surrogate approach, would show even greater advantage over traditional reduction methods.

Second, extending the framework to one-dimensional and multi-dimensional CFD simulations requires addressing two interconnected challenges. On the database generation side, training data must be representative of the broader thermochemical states encountered in reactive flows, including mixing layers, flame fronts and post-flame regions, which are absent from purely homogeneous 0D databases. Strategies such as curvature based oversampling or adaptive database enrichment from preliminary CFD runs should be explored to ensure adequate coverage of the chemical manifold. On the machine learning side, the network architecture and training procedure must be improved for ensuring robustness and stability of the NN surrogate in complex coupled simulations as it is a key prerequisite before deployment in realistic rocket engine configurations.

Acknowledgments

The authors would like to thank MaiaSpace and ONERA's scientific direction for funding and supporting the present work.

References

[1] Collectif Pour un réveil écologique, Report on space sector, https://pour-un-reveil-ecologique.org/documents/92/Rapport_complet.pdf (jan 2024).

[2] P. Pepiot-Desjardins, H. Pitsch, An efficient error-propagation-based reduction method for large chemical kinetic mechanisms, Combustion and Flame 154 (1) (2008) 67–81.

[3] A. Felden, E. Riber, B. Cuenot, Impact of direct integration of Analytically Reduced Chemistry in LES of a sooting swirled non-premixed combustor, Combustion and Flame 191 (2018) 270–286.

[4] P. O. Sturm, A. S. Wexler, A mass- and energy-conserving framework for using machine learning to speed computations: a photochemistry example, Geoscientific Model Development 13 (9) (2020) 4435–4442, publisher: Copernicus GmbH.

[5] F. M. Rohrhofer, S. Posch, C. Gößnitzer, J. M. García-Oliver, B. C. Geiger, Utilizing neural networks to supplant chemical kinetics tabulation through mass conservation and weighting of species depletion, Energy and AI 16 (100341) (May 2024).

[6] O. Owoyele, P. Pal, ChemNODE: A neural ordinary differential equations framework for efficient chemical kinetic solvers, Energy and AI 7 (100118) (Jan. 2022).

[7] S. Wu, W. Liang, K. H. Luo, Deep learning based combustion chemistry acceleration method for widely applicable NH_3/H_2 turbulent combustion simulations, Combustion and Flame 278 (114218) (Aug. 2025).

[8] T. Wang, Y. Yi, J. Yao, Z.-Q. J. Xu, T. Zhang, Z. Chen, Enforcing physical conservation in neural network surrogate models for complex chemical kinetics, Combustion and Flame 275 (114105) (May 2025).

[9] C. Mehl, D. Aubagnac-Karkar, 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 (067115) (Jun. 2023).

[10] X. Chen, Machine learning applied to the computation of chemical source terms in reacting flows, PhD thesis, Université Paris sciences et lettres (Jan. 2024).

[11] J. H. Halton, On the efficiency of certain quasi-random sequences of points in evaluating multi-dimensional integrals, Numerische Mathematik 2 (1) (1960) 84–90.

[12] K. He, X. Zhang, S. Ren, J. Sun, Deep residual learning for image recognition (2015).

[13] D. P. Kingma, J. Ba, Adam: A method for stochastic optimization (2014).

[14] S. Lloyd, Least squares quantization in PCM, IEEE Transactions on Information Theory 28 (2) (1982) 129–137.

[15] X. Chen, C. Mehl, T. Faney, F. Di Meglio, Clustering-Enhanced Deep Learning Method for Computation of Full Detailed Thermochemical States via Solver-Based Adaptive Sampling, Energy & Fuels 37 (18) (2023) 14222–14239, publisher: American Chemical Society.

[16] G. Strang, On the Construction and Comparison of Difference Schemes, SIAM Journal on Numerical Analysis 5 (3) (1968) 506–517, publisher: Society for Industrial and Applied Mathematics.

[17] L. F. Shampine, M. W. Reichelt, The MATLAB ODE Suite, SIAM Journal on Scientific Computing 18 (1) (1997) 1–22, publisher: Society for Industrial and Applied Mathematics.

[18] A. Matynia, A. Delaroque, H. Chakravarty, Brookesia: An open-source tool for reduction and optimization of kinetic mechanisms, https://github.com/Brookesia-py/Brookesia, version 1.8.4 (2021).

[19] A. Kronenburg, R. W. Bilger, J. H. Kent, Modeling soot formation in turbulent methane–air jet diffusion flames, Combustion and Flame 121 (1) (2000) 24–40.