ACS ENGINEERING Au
AN OPEN ACCESS JOURNAL OF THE AMERICAN CHEMICAL SOCIETY

pubs.acs.org/engineeringau
Article

Physically Consistent Neural Network Surrogates of Catalytic Reactors with Detailed Surface Kinetics

Published as part of ACS Engineering Au special issue "AI and Machine Learning in Chemical Engineering: Breakthroughs and Applications".

Felix A. Döppel, Mauro Bracconi,* and Matteo Maestri*

Cite This: ACS Eng. Au 2026, 6, 359–371
Read Online

ABSTRACT: Motivated by their potential to unlock real-time control of catalytic devices and model-based process design under full mechanistic detail, we developed physics-informed neural networks (PINNs) that consider detailed surface kinetics through an auxiliary kinetic surrogate model and enforce atom conservation as a hard constraint. In particular, we inform the PINNs about detailed surface kinetics by coupling with an auxiliary neural network that determines the surface reaction rates. This allows meshless predictions of composition profiles and removes the most time-consuming step from model training. Further, we propose a dedicated neural network layer that imposes atom conservation as a hard constraint. This guarantees the physical consistency of predicted reactor compositions and thus reduces the amount of data required for model training, in our case, by a factor of 10. We demonstrate that our new framework provides accurate and consistent concentration profiles of a CO_2 methanation reactor more than 1000 times faster than established direct integration schemes.

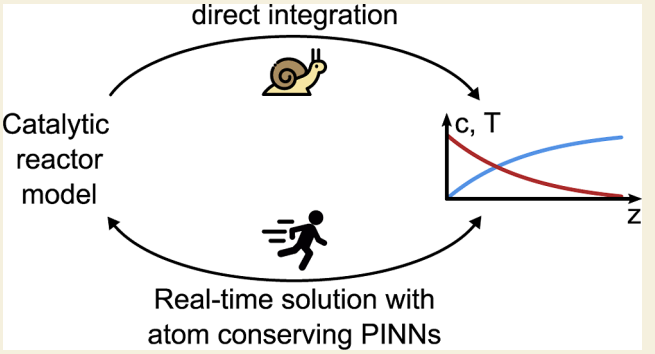

KEYWORDS: *physics-informed neural networks, catalysis, hard constraints, surrogate modeling, chemical kinetics, atom conservation*

INTRODUCTION

Detailed numerical modeling is essential for the optimization, intensification, control, and design of chemical and catalytic reactors. In particular, multiscale modeling has shown great potential due to its ability to capture physical phenomena across all relevant time and length scales. In this context, catalytic reactors are described by partial differential equations (PDEs), which are typically extremely stiff, making their solutions with classical methods computationally demanding. This is particularly true for heterogeneously catalyzed systems, where the reactivity is additionally affected by the surface morphology. Therefore, multiscale models come at the expense of high computational demand, often resulting in simulation times that are incompatible with time-sensitive applications, such as real-time control. While reduced-order models are oftentimes employed to reduce the simulation time, they also introduce simplifying assumptions, hindering their ability to describe complex phenomena. In contrast, machine learning-based surrogate models are computationally efficient by design and capture a high level of detail without being bound to theoretical assumptions. While they have been used to reduce the simulation time of industrial-scale catalytic reactor models from months down to days,[1] many engineering applications demand even faster, near-instantaneous responses. This can be achieved by replacing the entire simulation with a machine learning surrogate. In this case, the physical consistency of the surrogate is a major factor, as it increases interpretability, improves robustness toward imperfect and/or missing data, and guarantees physically meaningful predictions, even in extrapolation settings.[2] Therefore, we assess efficient and physically consistent surrogate models of catalytic reactors, considering detailed surface kinetic schemes.

Physics-informed neural networks (PINNs) have recently gained popularity as an effective tool for approximating the solution of PDEs. In contrast to standard neural networks, the objective function of PINNs also considers physical information. Automatic differentiation is used to conveniently compute the gradients of the objective function with respect to model parameters.[3] Once trained, PINNs provide a meshless

Received: November 12, 2025
Revised: January 8, 2026
Accepted: January 9, 2026
Published: January 30, 2026

https://doi.org/10.1021/acsengineeringau.5c00106
ACS Eng. Au 2026, 6, 359–371

and extremely efficient solution to PDEs, as demonstrated in the literature for numerous applications.[2,4] The meshless characteristic makes them particularly suitable for reactor model surrogates, as it enables evaluation at any position in the reactor at any reaction time. For example, Ngo and Lim used PINNs to predict concentration profiles of catalytic CO_2 methanation in an isothermal fixed-bed reactor, assuming a reactive system with five species and a single reaction following Langmuir−Hinshelwood−Hougen−Watson (LHHW)-type kinetics.[5] Lastrucci et al. described the profiles of composition, temperature, and pressure along a fixed-bed reactor for methanol synthesis using a reactive system of seven species and two reactions following LHHW-type kinetics.[6] In another example,[7] they additionally considered external heat transfer and predicted hotspot temperatures. While such simplified kinetic models are often sufficient to describe the overall reactor behavior, more detailed surface mechanisms are necessary to capture complex effects, such as morphology dependencies and changes of the rate determining step.[8] In this context, kinetics-informed neural networks, a variant of PINNs, have been developed to create surrogate approximators for the solution of zero-dimensional surface kinetic models[9,10] and to discover kinetic parameters from experimental data.[11] However, the combination of detailed surface kinetics with PINN-like surrogate models of one- or higher-dimensional catalytic reactors is currently missing.

Here, we close this gap by developing such PINN surrogates of catalytic reactors that specifically include detailed surface kinetics. A major obstacle for this is the difficulty in providing the optimizer with the correct gradient of the physical loss with respect to the model parameters. The catalyst surface typically equilibrates quasi-instantaneously with the fluid phase,[12] and as the steady-state surface composition θ is given by an implicit function $d\theta/dt = 0$ rather than an analytical expression, the required gradient is not directly accessible through automatic differentiation. To address this, we propose to describe the steady-state surface kinetics with a differentiable surrogate model that natively supports automatic differentiation instead of repeatedly solving the implicit equation during training. For this purpose, we propose to use neural network-based models since they have emerged as the most promising type of kinetic surrogate, overcoming the performance of splines,[13−16] Shepard interpolation,[17−19] and other machine learning-based surrogates like random forests[20,21] and kernel models.[22] The superior performance is linked to their intrinsic ability to structurally implement a priori physical knowledge such as the rate of the rate-determining step,[23] the Arrhenius law,[24,25] mass and atom conservation,[26−31] as well as thermodynamics.[26−28] This results in physically consistent steady-state source term predictions that provide accurate solutions up to 10,000 times faster than state-of-the-art ODE solvers.[23] The resulting expressions for fluid phase species source terms are compatible with automatic differentiation, effectively solving the challenge of gradient propagation.

However, a critical shortcoming of the PINN remains, namely, the nature of its soft constraints. Hard constraints are known to be more effective because they fully ensure the constraint to be met instead of just approximating it. In particular, atom conservation is a fundamental law that affects all reactive systems, including the catalytic reactors considered in this work. As it belongs to the class of linear equality constraints, it can be hard constrained by projection or completion.[32,33] In the projection approach, nonlearnable projection matrices are used to map preliminary model predictions onto the subspace of physically feasible solutions.[34] Frameworks, such as Hardnet,[33] embed a projection step into neural networks to enforce certain equality and inequality constraints. Another example is the ENFORCE method, which applies an iterative procedure to guarantee that neural network predictions satisfy any given nonlinear constraint within a specified tolerance ϵ.[32] In the context of chemical engineering applications, projection was recently used to close the atom balance in neural network predictions for process simulations,[35] and to close the nonlinear energy balance of catalytic reactors with LHHW-type kinetics.[7] In cases where projection is applied to physical quantities such as mass fractions and temperatures, it is important to preserve the positivity of the projected solutions. To this end, we will extend the approach by including a weighting matrix that adjusts the direction of the projection step, aiming at minimizing the possibility that negative values are obtained for physical quantities. In parallel to this work, similar approaches have been proposed, which aim to preserve positivity by performing orthogonal projection in a scaled space.[36,37] Differently, in the completion approach, only a subset of the output variables is predicted. The remaining outputs are calculated analytically to fulfill the constraints. It has been used in numerical models of the atmosphere[38] and physically consistent reaction pathway discovery.[31]

We will therefore include atom conservation hard constraints into PINNs and evaluate the effectiveness of completion and different (weighted) projection variants, aiming at demonstrating their critical importance for surrogate performances.

As a result, we present efficient and physically consistent neural network surrogates that replace the numerical solution of catalytic reactor models, considering detailed surface kinetics, as demonstrated here for a CO_2 methanation reactor model. This is made possible by two key advances. On one side, we enable PINNs to incorporate detailed surface kinetic schemes by replacing the implicit steady-state evaluations with physically consistent, differentiable surrogates. This overcomes the gradient propagation issue that has so far hindered model convergence. On the other side, we introduce an atom conservation hard constraint layer for PINNs that simultaneously avoids negativity and the numerical overhead related to condition-specific projection weights.

Overall, our contributions unlock a new class of PINN-based catalytic reactor surrogates that are both fast and physically consistent, paving the way for real-time, high-fidelity modeling of catalytic systems.

METHODS

Imposing Atom Conservation with a Dedicated Neural Network Layer

The fundamental law of atom conservation

$$F \cdot \omega = \text{const.} \tag{1}$$

with the atom-molecule matrix of mass fractions F (for example, Table S2 in the Supporting Information) and the mass fractions ω is a linear equality constraint acting on all chemically reacting systems. Specifically, the product $F \cdot \omega$ yields the mass fraction of each chemical element in the mixture. As the total mass and masses of each element in the mixture remain constant, the product $F \cdot \omega$ must be constant, as well.

Therefore, mass fractions inside a reactor predicted by any surrogate model $\hat{\omega}$ have to fulfill

https://doi.org/10.1021/acsengineeringau.5c00106
ACS Eng. Au 2026, 6, 359−371

$$F\cdot\omega_0 - F\cdot\hat{\omega} = 0 \tag{2}$$

with the mass fractions at the reactor inlet ω_0. This means that the predicted mass fractions conserve the number of elemental species present at the reactor inlet.

Recently, Chen et al.[34,35] proposed to enforce this constraint through orthogonal projection of $\hat{\omega}$ onto $\tilde{\omega}$ in the physically feasible solution space as

$$\tilde{\omega} = C\cdot\omega_0 + (I - C)\cdot\hat{\omega} \tag{3}$$

$$=\hat{\omega} + C\cdot(\omega_0 - \hat{\omega}) \tag{4}$$

with the unit matrix I and the correction matrix C (for example, eq S21 in the Supporting Information).

In doing so, the predicted mass fractions $\hat{\omega}$ are projected onto $\tilde{\omega}$, which lies within the physical subspace of mass fractions that conserve the amount of elemental species. The correction matrix C can be computed from F as

$$C = F^T(FF^T)^{-1}F \tag{5}$$

where F^T notates the transpose of F.

The atom conservation layer (Figure 1) implements the linear correction step (eq 3) in a single hidden layer with inputs $\hat{\omega}$, weights $(I - C)$, the bias $C\cdot\omega_0$, and no nonlinear activation function. As described within the results section, a generalized form of the matrix C (eq 10) might be employed to increase accuracy and preserve positivity of $\tilde{\omega}$.

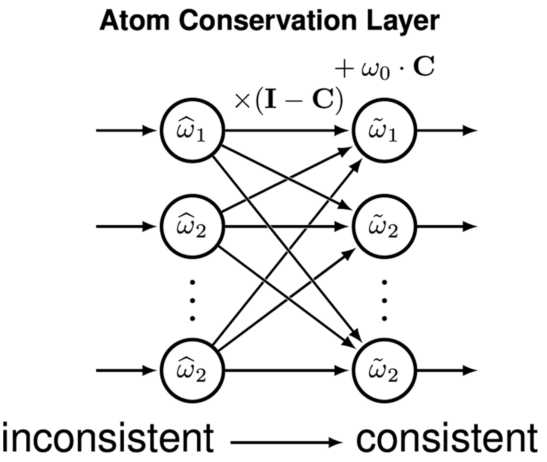

Figure 1. Schematic representation of the atom conservation layer. Predicted, typically atom balance violating (inconsistent) mass fractions $\hat{\omega}$ are mapped onto atom-conserving (consistent) mass fractions $\tilde{\omega}$ by multiplication with the matrix $(I - C)$ and subsequent addition of the product between the mass fractions in the feed ω_0 and C. No activation function is applied.

Atom-Conserving Neural Networks

We propose the atom-conserving neural network (AC-NN, Figure 2) to predict physically consistent reactor outlet compositions as a function of the reactor feed conditions, such as the feed temperature T_0 and the mass fractions ω_0 in the feed. AC-NNs consist of a fully connected neural network, for example, with hyperbolic tangent activation in the hidden layers. Exponential activation is used for the output layer to guarantee the positivity of predicted compositions and allow the model to effectively capture composition values over different orders of magnitude. After the classic output layer, we further add the atom conservation layer (Figure 1) to impose atom conservation as a hard constraint.

During training, the model minimizes the data loss $\mathcal{L}_{\text{data}}$

$$\mathcal{L}_{\text{data}} = \frac{1}{N_{\text{train}}}\cdot\sum_{n=1}^{N_{\text{train}}}\left(\text{asinh}\left(\frac{\omega_{i,n}}{\text{rref}_{\text{data}}}\right) - \text{asinh}\left(\frac{\tilde{\omega}_{i,n}}{\text{rref}_{\text{data}}}\right)\right)^2 \tag{6}$$

with the amount of training data N_{train}, the inverse hyperbolic sine function asinh, the true $\omega_{i,n}$ and estimated $\tilde{\omega}_{i,n}$ mass fraction of species i and sample n, and the reference value $\text{rref}_{\text{data}}$ of 10^{-6}. The reference value acts as a threshold below which the different orders of magnitude should not be resolved anymore because the respective values do not contribute significantly to the behavior of the system.[14,24] Here, we chose a threshold of around 1 ppm for the mass fractions. However, this parameter can be adapted depending on the problem of interest. The pseudologarithmic asinh is used as it approximates a relative error measure while avoiding numerical instabilities whenever any prediction is several orders of magnitude off.[24] Training is performed for 1000 epochs with the L-BFGS optimizer[39] using strong Wolfe line-search and an initial learning rate of 10^{-6}.

Physics-Informed Neural Networks for Catalytic Reactors with Atom Conservation Hard Constraints (AC-CatalyticReactor-PINN)

We further extend the previous approach to the prediction of full reactor profiles in terms of temperature and composition. Here, the resulting atom-conserving catalytic reactor PINN (AC-CatalyticReactor-PINN, Figure 3) is showcased for an isothermal plug flow reactor model. The extension of the proposed approach to more complex reactor models is straightforward.

In addition to the data loss considered already by the AC-NN, the AC-CatalyticReactor-PINN considers detailed surface kinetic schemes via a physics loss $\mathcal{L}_{\text{physics}}$, which is defined as

$$\mathcal{L}_{\text{physics}} = \frac{1}{N_{\text{physics}}}\cdot \sum_{n=1}^{N_{\text{physics}}}\left(\text{asinh}\left(\frac{\dot{s}(\tilde{\omega}_{i,n}, T)/\tilde{\omega}_{i,n}}{\text{rref}_{\text{physics}}}\right) - \text{asinh}\left(\frac{\dot{\tilde{\omega}}_{i,n}/\tilde{\omega}_{i,n}}{\text{rref}_{\text{physics}}}\right)\right)^2 \tag{7}$$

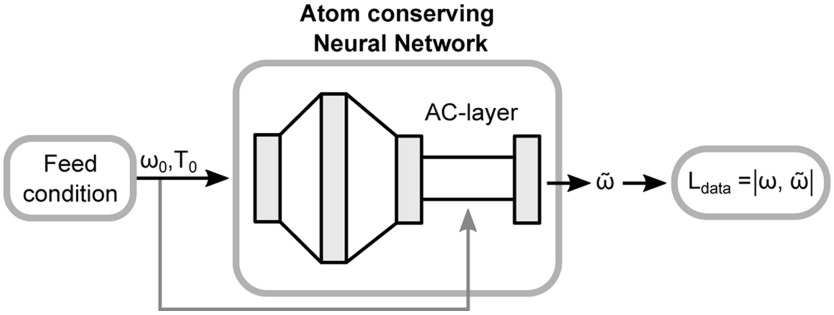

Figure 2. Schematic representation of the atom-conserving neural network (AC-NN). The atom conservation layer (AC-layer) imposes physically consistent predictions $\tilde{\omega}$. Its parameters depend on the feed condition.

https://doi.org/10.1021/acsengineeringau.5c00106
ACS Eng. Au 2026, 6, 359−371

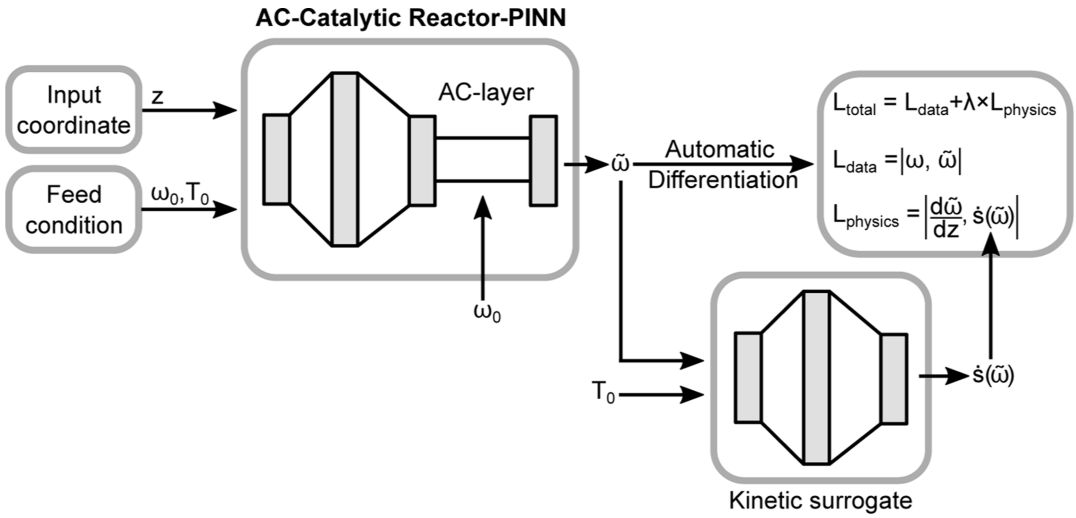

Figure 3. Schematic representation of the atom-conserving catalytic reactor PINN (AC-CatalyticReactor-PINN). The parameters of the atom conservation layer depend on the feed conditions. The kinetic surrogate supports automatic differentiation and efficiently provides physically consistent steady-state source terms of a detailed surface kinetic scheme as a function of the reaction conditions at the input coordinate, as predicted by the AC-CatalyticReactor-PINN.

with the number of collocation points $N_{physics}$, source terms $\dot{s}$, the derivative $\tilde{\omega}$ of estimated mass fractions with respect to the reactor coordinate, and the reference value $rref_{physics}$. The reference value acts as a threshold below which the different orders of magnitude should not be resolved anymore because the respective values do not contribute significantly to the behavior of the system.[14,24]

In the numerator, source terms and derivatives are divided by the predicted mass fractions to capture the relative change over many orders of magnitude instead of exclusively capturing the change of species with large mass fraction values. Here, we chose $rref_{physics}$ to be 5×10^{-2} m^{-1}. The pseudologarithmic asinh is further used to approximate a relative error measure while avoiding numerical instabilities whenever any prediction is several orders of magnitude off.[24] The evaluation of the physics loss includes the computation of steady-state source terms $\dot{s}$. This requires solving implicit functions $\frac{d\theta}{dt} = 0$ to find the steady-state surface coverages θ, i. e., those that are equilibrated to a fixed fluid phase composition, which is not directly compatible with automatic differentiation. Therefore, we replace the implicit computation of $\dot{s}$ with the call to a kinetic surrogate. The surrogate predicts the steady-state source terms as a function of the temperature and predicted gas phase mass fractions and is directly compatible with automatic differentiation.

The total loss $\mathcal{L}_{total}$ minimized by the optimizer during training is

$$\mathcal{L}_{total} = \mathcal{L}_{data} + \lambda \cdot \mathcal{L}_{physics} \tag{8}$$

where the commonly employed physical loss weight λ is set to 10^{α} for the AC-CatalyticReactor-PINN to balance the typically unequal contributions of physics and data. The AC-CatalyticReactor-PINN was trained for two different scenarios: in the first scenario, 40 different feed temperatures are sampled uniformly random from the range of 500 to 900 K and solved in the monolith reactor model with a fixed feed composition of 4% CO_2, 5.3% H_2, and 90.7% N_2, to supply the AC-CatalyticReactor-PINN with 40 inlet−outlet data pairs for consideration in the data loss during training. Additionally, 4000 collocation points are sampled uniformly random from the same temperature range and random axial reactor positions from the range of 0−1.2 cm. At those positions, the physics loss will be evaluated and minimized during training. The validation and test data sets are randomly sampled from the same ranges and solved for 200 and 10,000 conditions, respectively. In the second scenario, instead 160 inlet−outlet data pairs are sampled from the same feed temperature range and a varying feed composition of 4% CO_2, a uniformly random H_2/CO_2 ratio between 1.325 and 2.65 and N_2 as the remaining species. 8000 collocation points are sampled from the same temperature and H_2/CO_2 ratio range and at random positions between 0 and 1.2 cm. The validation and test data sets are randomly sampled from the same ranges and solved for 200 and 10,000 conditions, respectively. In both scenarios, α follows a linear ramp of −8 at epoch 15 to −2 at epoch 50. The final value of 10^{-2} is found empirically through preliminary work but can easily be replaced by adaptive loss weighting algorithms.[40] However, when testing a simple adaptive weighting procedure,[6] we observed worse performance or poor convergence (see Supporting Information). This is due to the strong sensitivity of the gradient loss toward the quality of the predictions. Instead, the weight of the gradient loss should be increased gradually at the beginning of the training phase. Training is performed for 1000 epochs with the L-BFGS optimizer[39] using strong Wolfe line-search and an initial learning rate of 10^{-6}.

Kinetic Surrogate

A kinetic surrogate model is employed to predict the steady-state source terms $\dot{s}_i$ of gas phase species as a function of the species mass fractions ω and the temperature T. As a case study, we consider a pseudohomogeneous reactor model representing the washcoated monolith reactor used in experiment 6 of Schmider et al.[41] (Section S1 in the Supporting Information). Accordingly, the underlying microkinetic scheme (Section S2 in the Supporting Information) is adopted from the same source. It consists of 42 elementary reactions among 19 species.[41] This kinetic mechanism is approximated by an ASINH-Net-like neural network surrogate,[24] which is fully differentiable and compatible with automatic differentiation frameworks. The inputs (ω, T) are converted to the logarithm of mass fractions and the inverse temperature before applying a scaler that maps the range of converted training data inputs to the interval (−1, 1). Further, the model has three hidden layers with 20 nodes each, resulting in 1042 adjustable model parameters in total. The output layer has two nodes with exponential activation corresponding to the source terms of the so-called key species. Here, we chose methane and hydrogen as the key species. In analogy to our previous work considering stoichiometric coefficients,[31] their source terms are then multiplied with the atom conservation matrix (Table 1), which is obtained as the reduced row echelon form of the null space of the atom-molecule matrix of mass fractions F (Table S2 in the Supporting Information). This operation can be seen as a completion step and yields the steady-state source terms for all gas phase species in the system.[42] It is important to note that this approach does not rely on the specific architecture of the ASINH-Net. In fact, many other types of surrogate models are suitable as well, including the formation−

https://doi.org/10.1021/acsengineeringau.5c00106
ACS Eng. Au 2026, 6, 359−371

Table 1. Atom Conservation Matrix for CO_2 Methanation Assuming CH_4 and H_2 as the Key Species in Analogy to Refs 31 and 42

	CH_4	H_2
CH_4	1	0
H_2	0	1
H_2O	−9/4	−9
CO_2	33/4	22
CO	−7	−14
N_2	0	0

consumption neural network,[42] the global reaction neural network,[26] or the kinetics-constrained neural network.[28] They have been proven effective for a wide range of catalytic systems under industrially relevant conditions, including the preferential oxidation of CO,[23] ammonia oxidation,[24] steam reforming,[1] direct DME synthesis,[42] and CO_2 Fischer−Tropsch.[28] Also, noncatalytic homogeneous reaction mechanisms like GRI3.0 have been accurately described,[42] indicating scalability of the surrogate-assisted training approach to systems with hundreds of species and reactions.

Training, validation, and test data sets for the generation of the kinetic surrogate were generated by solving the monolith reactor model (Section S1 in the Supporting Information) at 200 equally spaced temperatures within 400−1000 K at a molar inlet composition of 4% CO_2, 5.3% H_2, and 90.7% N_2, resulting in 149,560 pairs of reaction conditions and corresponding source terms. The amount of data was downsampled by saving only every 10th pair, and the remaining 14,956 pairs were randomly split to 7470 data for training, 3743 for validation, and 3743 for testing. The data sets for the kinetic surrogate used for the scenario with different feed compositions were generated in the same way, but with a molar inlet composition of 4% CO_2 and 5.3% to 10.6% H_2 in N_2. The resulting trajectories were randomly split into 7470 data for training, 3743 for validation, and 3743 for testing.

The L-BFGS optimizer[39] with strong Wolfe line-search is employed to reduce the thresholded relative error in the source terms[14,23]

$$e_{\mathrm{rel}}^{\mathrm{threshold}} = \left| \frac{\dot{s} - \dot{s}_{\mathrm{pred}}}{\max(\dot{s}, \mathrm{threshold})} \right| \tag{9}$$

in the training data set until the validation error did not improve for 1000 consecutive epochs. The threshold is chosen as 10^{-8} m^{-1}. For the final model, the weights corresponding to the lowest validation error are chosen. The test error at unseen conditions is 0.25% (Figure S1 in the Supporting Information).

RESULTS AND DISCUSSION

We propose the atom-conserving catalytic reactor PINN (AC-CatalyticReactor-PINN) as a physically consistent and numerically efficient surrogate to catalytic reactor models based on a detailed surface kinetic mechanism.

To ensure the physical consistency of the model predictions, we first study suitable methods of imposing hard constraints on atom conservation on reactor surrogates in general. We embed these methods into feed-forward neural networks, obtaining the atom-conserving neural network (AC-NN), and use it as a surrogate to the outlet composition of a CO_2 methanation reactor. Further, we evaluate the robustness of these methods toward noise using synthetic reactor data. We find that the orthogonal projection amplifies noise and can lead to negative predictions whenever species with small concentrations are present. In contrast, weighted projection and completion both impose atom conservation, increase the accuracy of predictions, and preserve the positivity of the initial solutions. However, we find that completion is more robust toward noise, adds a significantly smaller computational overhead to the model evaluation, and requires fewer training data than the weighted projection approach.

Finally, we inform the AC-NN about detailed surface kinetics under steady state. In particular, we employ previously developed physically consistent surrogates of the surface kinetic scheme to enable convenient gradient propagation through automatic differentiation and simultaneously reduce the training time of the model. We apply the resulting AC-CatalyticReactor-PINN to build surrogates of a CO_2 methanation reactor considering the detailed surface kinetic mechanism proposed by Schmider et al.[41] This yields accurate (<4% relative error) and physically consistent predictions of the composition trajectories at a 1000 times lower numerical cost than Cantera. Cantera is an open-source chemical kinetics software used for solving chemically reacting laminar flows.[43]

Imposing Atom Conservation on Reactor Surrogates

Atom conservation (eq 1) is a fundamental physical law acting on all chemically reactive systems. Figure 4 schematically shows different possibilities of imposing it to neural network predictions. For example, Chen et al. proposed a linear correction step,[34,35] which can be understood as the orthogonal projection of any predicted composition (blue symbol in Figure 4) onto the closest point in the atom-conserving manifold (red symbol in Figure 4). However, this procedure has been shown to create physically infeasible compositions, such as negative concentrations,[37] which might accumulate and destabilize reactive simulations.[36,44] Here, we explore weighted projection (yellow symbol in Figure 4) and completion (green symbol in Figure 4) as alternatives to the orthogonal projection approach and compare them regarding accuracy, positivity, and computational overhead. First, the methods will be presented in detail and applied to impose atom conservation to neural network surrogates of a CO_2 methanation reactor, predicting its outlet compositions. Then, we will compare the methods regarding their robustness toward noise.

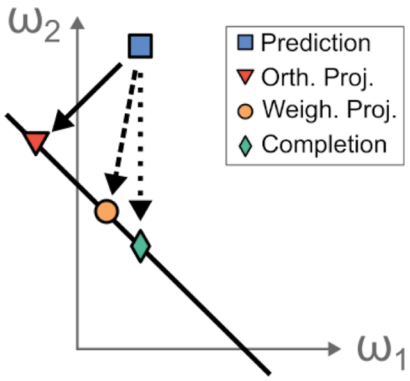

Figure 4. Schematic representation of atom conservation hard constraints acting on predicted compositions. In a system composed of two components, 1 and 2, predicted mass fractions ω (blue square) are likely to violate the atom balance (solid line). Orthogonal projection yields the closest atom-conserving composition (red triangle), which may, however, contain unphysical negative values. Weighted projection (yellow circle) corrects the smaller values less than the larger ones (here: $\omega_1 < \omega_2$) in order to preserve positivity. The completion step (green diamond) avoids predicting ω_2 and instead assigns it a physically valid value based on the other component(s).

https://doi.org/10.1021/acsengineeringau.5c00106
ACS Eng. Au 2026, 6, 359−371

Application to Methanation Reactor Surrogates

Here, we compare the different methods of imposing atom conservation to neural network surrogates, using the example of a catalytic CO_2 methanation reactor. This case study involves an isothermal 1-D pseudohomogeneous reactor model with fixed reactor geometry, catalyst mass, and total gas flow. It corresponds to the experimental study N° 6 by Schmider et al.[41] A feed-forward neural network is employed to predict the outlet composition as a function of the temperature, assuming a stoichiometric feed. Training data (symbols) are drawn from simulated outlet compositions (full lines) at seven equally spaced temperatures in the range of 550 to 850 K (Figure 5).

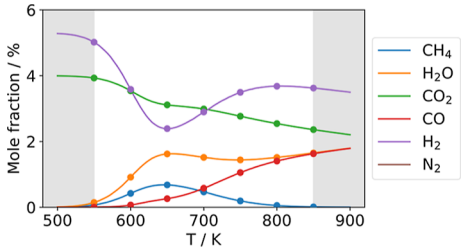

Figure 5. CO_2 methanation reactor outlet composition as a function of the feed temperature. Model training data (symbols) are equidistantly sampled from the exact solutions (solid lines) at seven different temperatures in the range 550 to 850 K. The gray region outside this range is considered extrapolation.

Predicting the outlet compositions with a standard neural network yields low agreement with the underlying ground truth (Figure 6). The data loss for unseen test data is 2.4×10^{-3} in interpolation and 4.0×10^{-1} in extrapolation. Further, the results violate the atom balance with a value of up to 10%.

Orthogonal Projection. In order to fix the physical inconsistency of the standard neural network model, atom conservation is imposed by using the orthogonal projection method. To this end, the atom conservation layer is appended to the standard neural network. This layer uses a set of precomputed weights and biases, which are obtained from the atom-molecule matrix of mass fractions F. While the orthogonal projection successfully eliminates any error in the atom balance of the predicted compositions, it has some structural shortcomings that become evident whenever species with small mole fractions are present. For example, the validation loss for CH_4 mole fractions increased by a factor of 11, and the overall test loss increased by a factor of 3.5 in interpolation and 70 in extrapolation. Further, the model predicts negative mole fractions for 42% of the test conditions in extrapolation (Figure 7). This is explained by the fact that the corrections resulting from orthogonal projection are of equal magnitude for all species, irrespective of their mole fraction values. Consequently, the corrections are often larger than the predicted values themselves, leading to high relative errors and, in extreme cases, even negative predictions.

Weighted Projection. To improve the handling of species with small mass fractions and simultaneously avoid negativity, we propose to change the direction of the projection step by applying a weight matrix W in the computation of the correction matrix C as

$$C_{\text{weighted}} = (WW^T)^{-1}F^T(F(WW^T)^{-1}F^T)^{-1}F \tag{10}$$

For weights W equal to the identity matrix I, the weighted projection is identical to the orthogonal projection and $C_{\text{weighted}} = C$.

Inspired by the recent work of Sturm and Silva,[45] we consider $W = \text{diag}(\hat{\omega}^{-1})$ to change the projection direction such that small mass fraction values are corrected less than

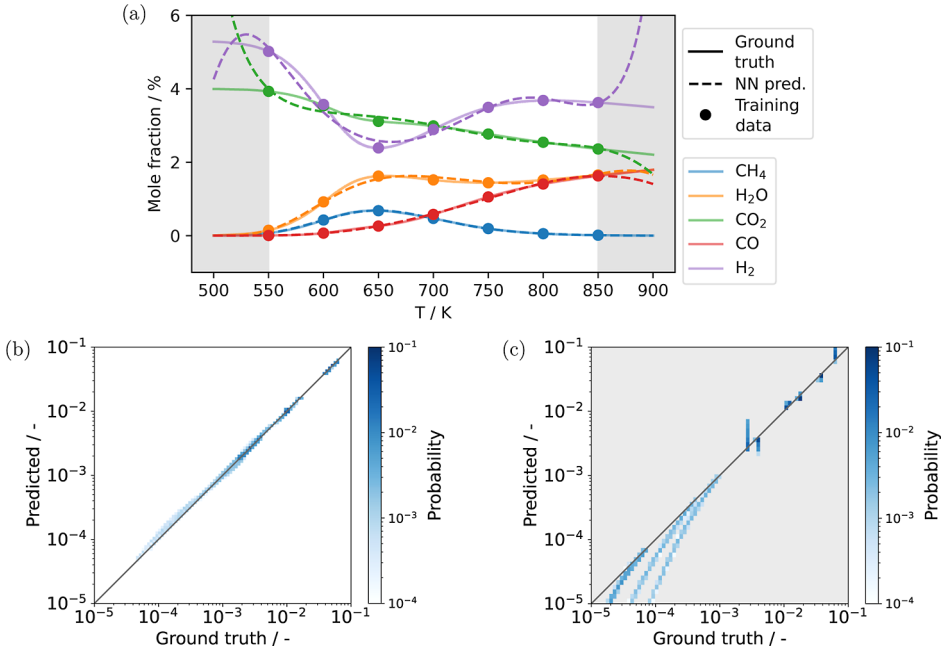

Figure 6. Methanation reactor outlet compositions predicted by a standard neural network without imposed atom conservation (a). Parity (jPDF) plots comparing the predicted and exact mass fractions for test data in interpolation (b) and extrapolation (c). Those data are sampled with temperatures in the range of 550−850 K for interpolation, as well as 500 to 550 K and 850 to 900 K for extrapolation. The color scale denotes the joint probability density of the data pairs.

https://doi.org/10.1021/acsengineeringau.5c00106
ACS Eng. Au 2026, 6, 359−371

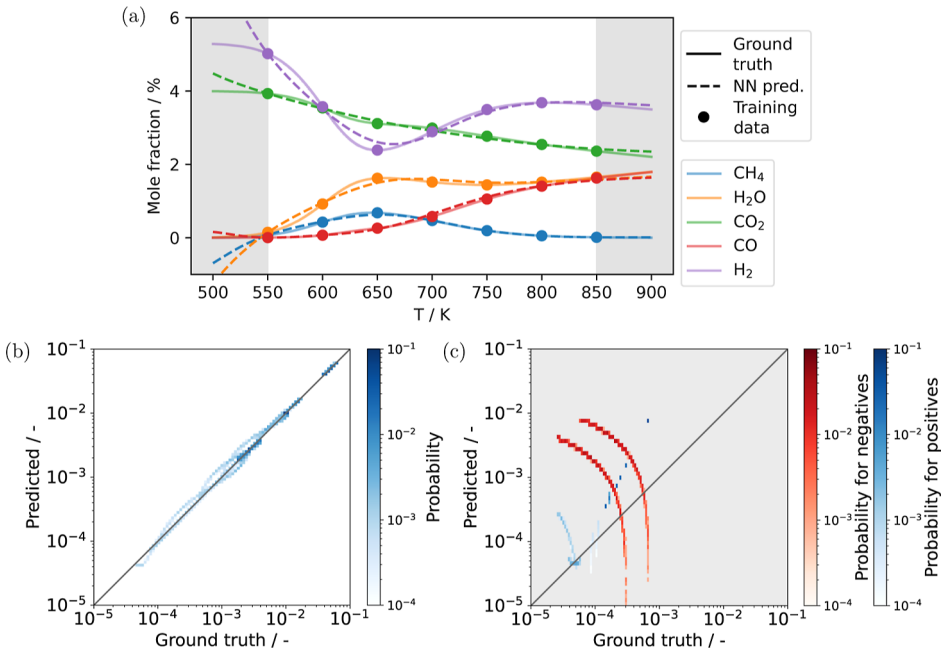

Figure 7. Methanation reactor outlet compositions predicted by a neural network with orthogonal projection to impose atom conservation (a). Parity (jPDF) plots comparing the predicted and exact mass fractions for test data in interpolation (b) and extrapolation (c). Those data are sampled with temperatures in the range of 550 to 850 K for interpolation, as well as 500 to 550 K and 850 to 900 K for extrapolation. The color scale denotes the joint probability density of the data pairs. Predictions of positive mass fractions are shown in blue, while predictions of negative, and therefore physically meaningless, mass fractions are shown in red.

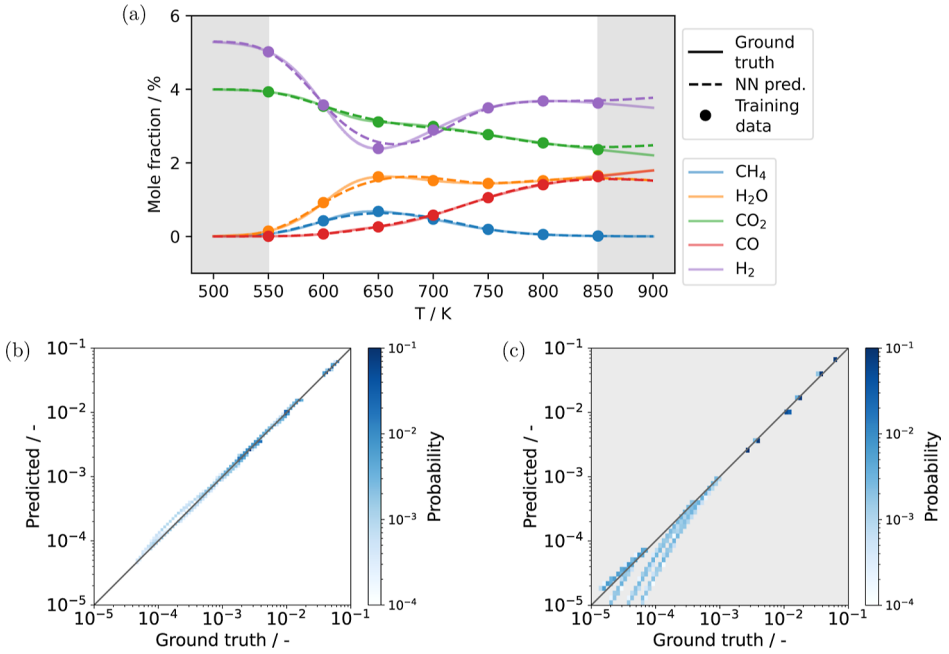

Figure 8. Methanation reactor outlet compositions predicted by a neural network with weighted projection to impose atom conservation (a). Parity (jPDF) plots comparing the predicted and exact mass fractions for test data in interpolation (b) and extrapolation (c). Those data are sampled with temperatures in the range of 550−850 K for interpolation, as well as 500 to 550 K and 850 to 900 K for extrapolation. The color scale denotes the joint probability density of the data pairs.

large ones (Figure 4). We will refer to this approach as weighted projection, although technically infinitely many choices of weights exist (Section S4 in the Supporting Information). In parallel to this work, similar approaches have been proposed,[36,37] which apply the projection to scaled versions of ω.

When the parameters of the atom conservation layer are chosen according to the weighted projection approach, predictions are much more accurate in the case of orthogonal projection. The test loss is reduced to 2.4×10^{-3} in interpolation and 2.9×10^{-1} in extrapolation. Further, no negative predictions occur (Figure 8), and the atom balance is perfectly closed.

https://doi.org/10.1021/acsengineeringau.5c00106
ACS Eng. Au 2026, 6, 359−371

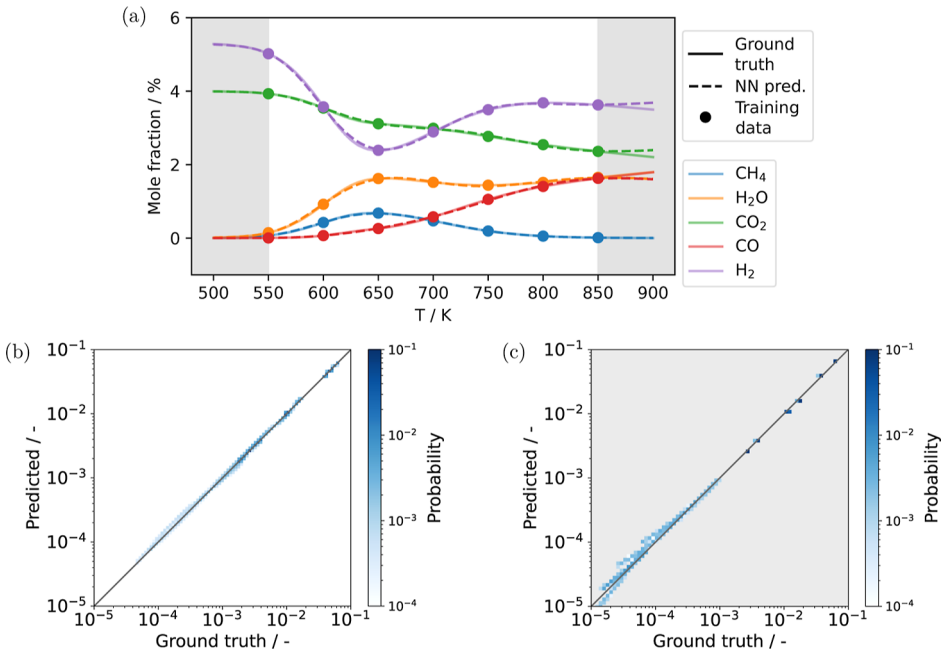

Figure 9. Methanation reactor outlet compositions predicted by a neural network with completion to impose atom conservation (a). Parity (jPDF) plots comparing the predicted and exact mass fractions for test data in interpolation (b) and extrapolation (c). Those data are sampled with temperatures in the range of 550−850 K for interpolation, as well as 500 to 550 K and 850 to 900 K for extrapolation. The color scale denotes the joint probability density of the data pairs.

Completion. In alternative to the instance-dependent weights of the weighted projection approach, we propose another set of weights given by the diagonal matrix

$$W_{i,i} = \begin{cases} 1 & \text{if } i \text{ is key species,} \\ 0 & \text{if } i \text{ is dependent species.} \end{cases} \tag{11}$$

where a subset of so-called key-species are assigned weights of one and all other dependent species are assigned a weight of zero.[1] Consequently, $\tilde{\omega}$ will not be influenced by the initial predictions of the dependent species mass fractions (Figure 4). To remove the unused parameters from the neural network model, we exclude the dependent species from the neural network output layer and the corresponding columns from the matrix $(I - C)$ (Section S5 in the Supporting Information). Interestingly, this reduced version of $(I - C)$ is identical to the reduced row echelon form of the null space basis of F.[2] This can be exploited to facilitate the implementation of eq 10, increase the performance, and avoid numerical issues arising from using zero weights.

In our previous work,[31] we used the reduced row echelon form of the null space basis of the mass fraction matrix as a completion matrix to enforce the physical consistency of stoichiometric coefficients discovered with a chemical reaction neural network. This demonstrates, that the completion method proposed by Beucler et al.[38] is formally a special case of projection. We will therefore refer to this approach as completion.

In contrast to the other projection approaches, the completion approach does not apply changes to the molar fractions of all species but rather preserves the predictions of the key species and completes the solution with physically consistent values for all the other, dependent species.

The resulting predictions have a test error of 8.4×10^{-4} for interpolation, which is three times more accurate than the best of the other methods, and 1.8×10^{-2} for extrapolation, which is about ten times more accurate than the other methods. No negative predictions occur, and the atom balance is respected (Figure 9).

One of the main practical differences between the weighted projection and the completion approach is that the completion uses a precomputed correction matrix, while with weighted projection, the matrix has to be computed individually for every single prediction. Therefore, in our case, the training time of the weighted projection model is five times higher than that of the completion model. A similar factor is expected for inference, but the exact value depends on the reaction conditions, the number of parameters, and the accuracy of the model, as well as the number of species in the reactive system.

Effect on the Required Amount of Training Data

Previously, we compared the atom conservation methods in terms of the accuracy obtained from a fixed amount of training data. Here, we compare them regarding the amount of training data required to reach a prediction accuracy of 1%. To this end, we tracked the relative prediction error of nine differently initialized neural networks for each of the four previously discussed methods and report the median result, i.e., the value for the fifth-best model (Figure 10). While the models without correction, with orthogonal projection, and with weighted projection yield overall comparable accuracies, the completion method yields the lowest relative error for any amount of training data. In particular, it requires only nine training points to yield a relative error of 1% on the validation data set. With orthogonal projection, this accuracy is reached with approximately 100 training data. Without correction and with the weighted projection, the target accuracy is never obtained within the considered range. While the exact results strongly depend on the complexity of the modeling task at hand, these

https://doi.org/10.1021/acsengineeringau.5c00106
ACS Eng. Au 2026, 6, 359−371

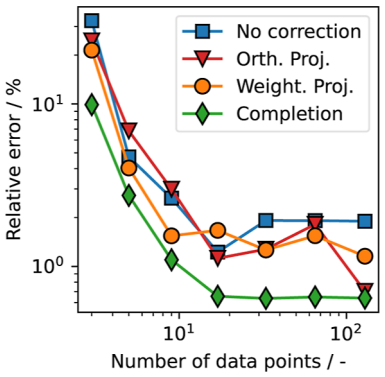

Figure 10. Prediction accuracy of neural network surrogates of methanation reactor outlet compositions for four different methods of imposing atom conservation. Results are shown as a function of the amount of training data used. The reported values are the median of the mean relative error obtained from nine differently initialized models each.

results indicate that the completion approach leads to a significant reduction in the required amount of training data.

Robustness toward Noise

Here, we apply the atom conservation methods to a data set of simulated methanation reactor outlet compositions, which are subject to multiplicative noise. These compositions are identical to the values shown in Figure 5 and were obtained by solving the methanation reactor model for feed temperatures between 500 and 900 K. The noise level in the perturbed data $\bar{\omega}$ is controlled by γ, which defines the bounds of the uniform distribution $\mathcal{U}(-\gamma, \gamma)$ as

$$\bar{\omega} = \omega \cdot (1 + \alpha), \text{ with } \alpha \sim \mathcal{U}(-\gamma, \gamma) \tag{12}$$

this noise introduces a relative error within $\pm\gamma$ mimicking inaccuracies arising from conventional neural network predictions at various stages of the training process and leads to a violation of the atom balance. The effect of atom conservation on such inaccuracies is central to the convergence behavior of a neural network in the training process. Alternatively, the noise can be seen to represent uncertainties from the experimental measurements.

Orthogonal Projection. Applying the orthogonal projections to the perturbed data aims to remove the atom conservation violations. However, by changing the composition values, it also affects the accuracy with respect to the original, unperturbed data. Naively, one might expect that the accuracy would increase as the unphysical part of the perturbation vanishes. However, the projection step actually lowers the accuracy, as indicated by the increase (Figure 11a) of the mean relative error

$$\epsilon_{\text{rel}} = \frac{1}{N_i}\frac{1}{N_j}\sum_i \sum_j \left| \frac{\bar{\omega}_{i,j} - \omega_{i,j}}{\omega_{i,j}} \right| \tag{13}$$

where N_i is the number of samples, N_j is the number of chemical species, $\bar{\omega}$ is the perturbed and potentially projected mass fractions, and ω is the original mass fractions.

This happens because the orthogonal projection step applies corrections of similar order of magnitude to all species, irrespective of the magnitude of their mass fractions. Therefore, whenever species with small mass fraction values are present, large relative corrections occur and increase the overall relative error. In extreme cases, the correction applied to small mass fraction values might even be larger than the original value. Then, the new composition can contain negative values. For this specific case study, this effect is observed even for small perturbations of $\gamma \leq 1\%$ (Figure 11b).

Weighted Projection. Not only does the weighted projection reduce the relative error of the perturbed data (Figure 11a), but it also avoids negative composition values up to a noise level of $\gamma \approx 70\%$ (Figure 11b). However, the adaptive nature of the weighted projection also presents a potential drawback. The correction matrix is dependent on the predicted mass fractions and must therefore be computed individually for each sample. This results in a numerical overhead approximately 3 orders of magnitude higher than that of the orthogonal projection, which uses a fixed, precomputed correction matrix (Figure 11c).

Completion. When applied to the perturbed composition data, the completion approach shows improved relative errors (Figure 11a), and preserves positivity for noise levels of up to $\gamma \approx 100\%$, indicating it is even more robust to noise than the weighted projection (Figure 11b). Importantly, the correction matrix used in the completion step is fixed, resulting in a minimal computational overhead comparable to that of the orthogonal projection (Figure 11c).

The effectiveness of the completion approach, however, is closely tied to the selection of key species. Consistent with our previous findings,[31] we show that species with the lowest overall concentrations should be selected as key species to guarantee the robustness of the approach. Indeed, choosing a high-concentration species as the key leads to large absolute corrections being applied to species with low concentrations, which amplifies relative errors. Low relative errors, in turn, are

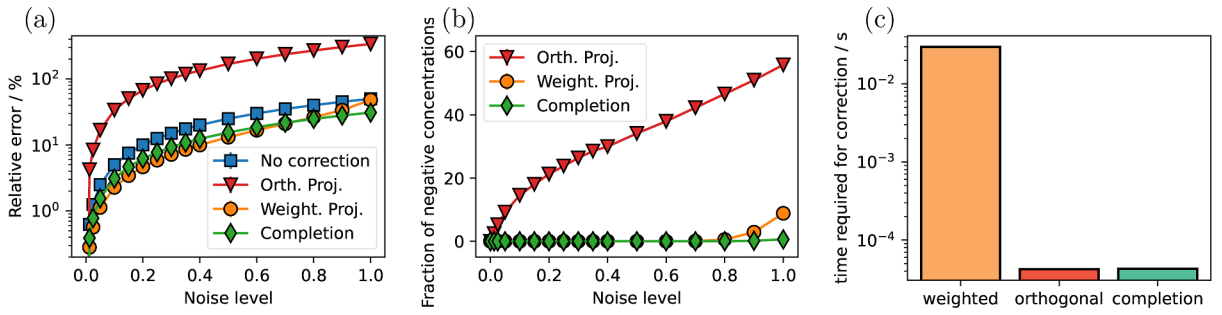

Figure 11. (a) Relative error between the corrected noisy compositions and the ground truth. (b) Fraction of corrected noisy compositions containing negative values. (c) Processing time of the different correction methods averaged over 50 runs and all noise levels. For (a,b), error bars are given by the standard deviation over 50 runs with different random seeds but are smaller than the symbol size and therefore not visible.

https://doi.org/10.1021/acsengineeringau.5c00106
ACS Eng. Au 2026, 6, 359−371

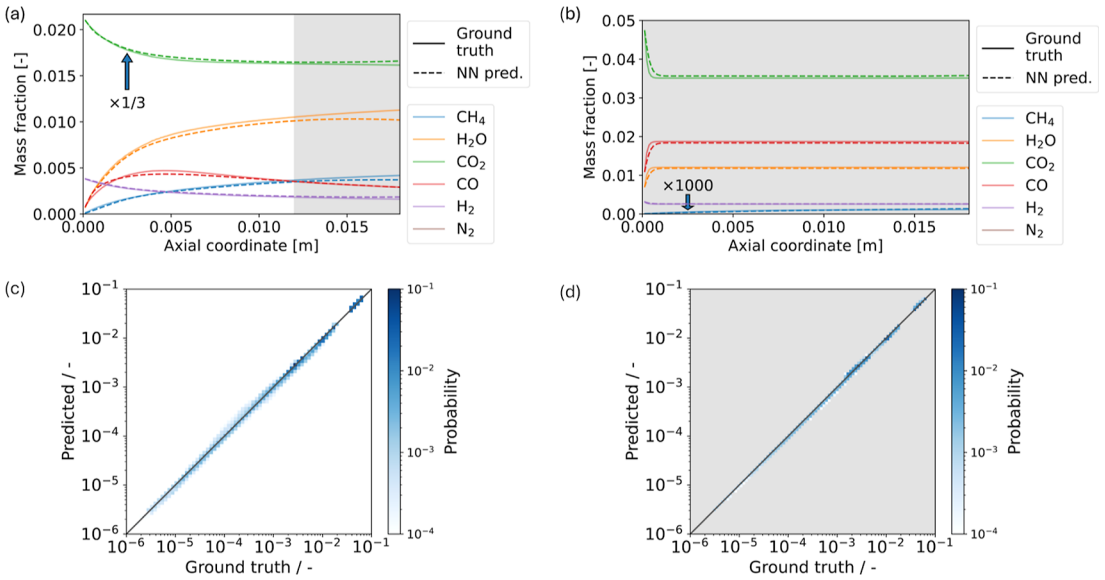

Figure 12. Reactor composition trajectory predicted at unseen conditions by the completion PINN (dashed) compared with the ground truth (Cantera solution, solid lines) (a) for T = 655 K (within training range) and (b) for T = 910 K (outside of training range). The assumed reactor length L is 1.2 cm. Data at later axial coordinates are shown to indicate extrapolation capabilities. Parity (jPDF) plots between the predicted and exact reactor trajectory mass fractions in the validation data set (c) for interpolation in z and (d) for extrapolation in z. Interpolation data are sampled with reactor positions z between 0 and L = 1.2 cm. Extrapolation data are sampled between L and 1.5·L. The color scale denotes the joint probability density of the data pairs.

essential for maintaining positive concentrations and ensuring the effectiveness of the physics loss via the species source terms. Section S6 in the Supporting Information shows this effect for the case of choosing CO_2, H_2O, or H_2 as the key species.

However, the choice of the key species only concerns the design of the completion matrix and does not limit practical applicability, since, independent of the choice, the model will predict mass fractions for all species of interest. This allows training our model even in the extreme case in which no information is available for the key species at all, simply by minimizing the data loss for the known (dependent) species. This is particularly important for training the model directly with the experimental data.

So far, we have tested the robustness of orthogonal projection, weighted projection, and completion toward synthetic, uniformly distributed noise. For completion, robustness has also been shown toward more realistic, systematic biases like missing data and sensor offsets during the automated discovery of kinetic mechanisms.[31] In future work, we will validate whether the findings of this work will hold for surrogates based on experimental data subjected to real-world measurement errors.

Coupling with Kinetic Surrogates to Enable Meshless Reactor Simulations

We now tackle the more complex modeling task of predicting the full composition profile in a methanation reactor as a function of the inlet temperature and axial position. We employ a PINN trained solely on inlet and outlet composition data. Physical information is accounted for by enforcing the governing equations at collocation points along the reactor, allowing accurate, meshless predictions at any position. For catalytic reactors, the governing equations contain the source terms resulting from the surface kinetic mechanism. Because the catalyst surface typically equilibrates quasi-instantaneously,[12] source terms must be evaluated with steady-state surface coverages. These coverages are defined by an implicit equation, which is usually solved by integrating a system of stiff differential equations. However, standard solvers for this task cannot be differentiated by using automatic differentiation. As a result, PINN training suffers from incomplete gradient information and poor convergence.

There are multiple approaches that would allow this challenge to be overcome, including differentiable solver implementations and implicit differentiation techniques. Among those, we identified the introduction of differentiable neural network surrogates of the steady-state surface kinetics as the most suitable approach because it evaluates 3 to 5 orders of magnitude faster than any alternative involving direct integration,[23,24,26] accelerating the training of the PINN accordingly. A detailed comparison of the different approaches is given in Section S7 of the Supporting Information.

The proposed kinetic surrogate predicts steady-state source terms directly from the reaction conditions. This enables full gradient propagation during training while remaining physically consistent, accurate, and computationally efficient. We pair this approach with the atom conservation hard constraints discussed above and refer to the resulting model as the atom-conserving catalytic reactor physics-informed neural network (AC-CatalyticReactor-PINN).

Application to Methanation Reactor Modeling with Changing Feed Temperatures

We evaluate AC-CatalyticReactor-PINN as a surrogate to a methanation reactor model with detailed surface kinetics, focusing on the impact of different strategies to impose atom conservation. Our findings are consistent with the observations obtained for simpler reactor models: omitting the atom conservation layer leads to physically inconsistent results. Compared to completion, it increases the validation loss by a factor of 4 for interpolation in the axial direction and a factor of 10 for extrapolation. The orthogonal projection approach maintains the atom balance but predicts negative concen-

https://doi.org/10.1021/acsengineeringau.5c00106
ACS Eng. Au 2026, 6, 359−371

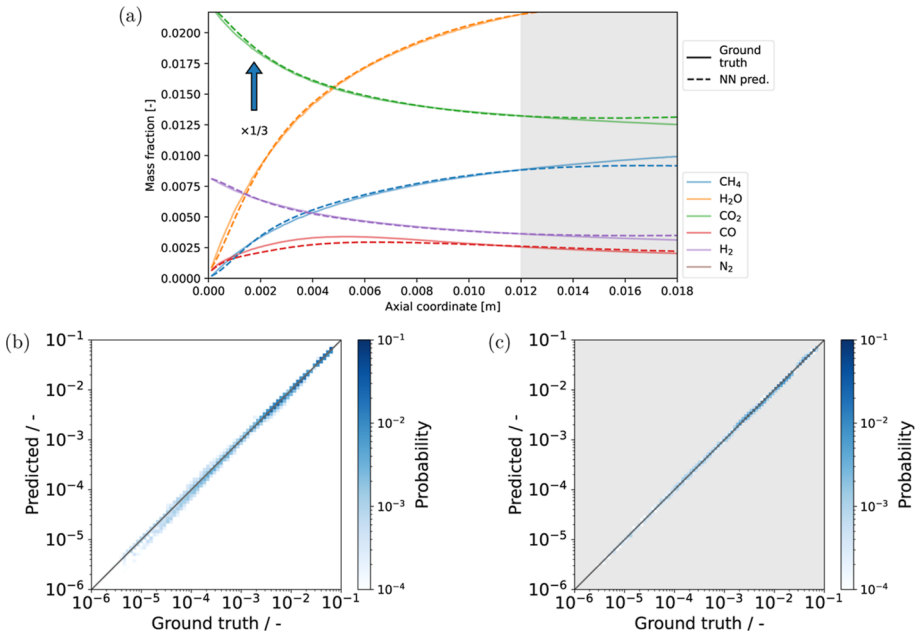

Figure 13. Compositions predicted by the completion, AC-CatalyticReactor-PINN based for different feed temperatures and compositions. (a) Reactor composition trajectory predicted at a unseen feed conditions (dashed) compared with the ground truth (Cantera solution, solid lines) for T = 655 K (within training range) and a feed composition ratio H_2/CO_2 of 2.65. The assumed reactor length is 1.2 cm. Data at later axial coordinates (gray area) are shown to indicate extrapolation capabilities. (b) Parity (jPDF) plot comparing the predicted and exact mass fractions trajectories in the validation data set with varying feed compositions. Results are shown for interpolation in the axial coordinate. The color scale denotes the joint probability density of the data pairs. Interpolation data are sampled with axial coordinates between L = 0 and 1.2 cm. (c) Parity (jPDF) plot for extrapolation in the axial coordinate sampled between L = 1.2 and 1.8 cm.

trations for 1% of the interpolation conditions and 21.6% of the extrapolation conditions. Such negative values are not only physically implausible but also incompatible with the kinetic surrogate, requiring the use of cutoff values that, in turn, undermine model convergence. The weighted projection approach successfully preserves positivity. However, its effectiveness depends on the prediction accuracy during training. In the early stages, inaccurate predictions lead to impractical projection directions, potentially amplifying the deviation from the ground truth. This effect is already evident in our method validation study with simulated composition data, where high noise levels lead to rapidly increasing relative errors (Figure 11a) and negative concentrations (Figure 11b). To mitigate this effect, we activated the weighted projection AC-layer only after 50 training epochs. However, this introduces a new training hyperparameter and does not fully resolve the convergence issue. In contrast, the completion approach is less sensitive to inaccurate predictions and therefore improves convergence from the very beginning of the training procedure. Based on these observations, we further consider only the completion-based AC-CatalyticReactor-PINN.

The completion-based AC-CatalyticReactor-PINN accurately captures the spatial concentration profiles of all species along the reactor (white area in Figure 12a) even for previously unseen temperatures from the wide range of feed temperatures covered (500−900 K). Even for extrapolation in the axial position (gray area in Figure 12a), the model provides physically plausible and quantitatively accurate results. For temperatures exceeding the training range, the reaction system becomes equilibrium-limited with most conversion occurring near the reactor inlet. Only minor amounts of methane are formed further downstream. Still, AC-CatalyticReactor-PINN extrapolates reliably, reproducing the correct behavior even for axial positions beyond the training range (Figure 12b). To assess the performance more generally, we evaluate the model on 100 random reaction trajectories within the temperature training range. The relative prediction error is 3.8% for interpolation (Figure 12c) and 2.8% for extrapolation in axial direction (Figure 12d). The lower extrapolation error is likely due to the flatter concentration profiles caused by lower reactivity near the reactor outlet. Atom balance is strictly enforced with negligible deviations below 10^{-6}%.

These results are particularly promising given the compact training setup: Only 40 inlet−outlet composition pairs and 4000 collocation points were used. This amount of compositions is experimentally accessible, and the low amount of collocation points enables model training within minutes on standard consumer hardware. With fewer than 1000 trainable parameters, the model remains lightweight and suitable for embedded or resource-constrained deployment scenarios. Another important advantage of using the AC-CatalyticReactor-PINN is that, in contrast to solutions obtained via direct integration, the provided solutions are meshless. This means that compositions can be obtained at any position within the reactor, instead of relying on spatial discretization schemes.

Application to Methanation Reactor Modeling with Changing Feed Temperatures and Compositions

To further demonstrate the potential of the AC-CatalyticReactor-PINN, we extend our previous experiments of modeling the composition profiles within the feed temperature range of 500−900 K by additionally varying the ratio of H_2 to CO_2 in the feed from 1.325 to 2.65. To this end, we collected 160 inlet−outlet composition pairs sampled from a regular grid within this range of feed conditions. The number of collocation points was increased to 8000, randomly sampled within the

https://doi.org/10.1021/acsengineeringau.5c00106
ACS Eng. Au 2026, 6, 359−371

range of feed conditions and axial positions, to cover the increased range of reaction conditions.

As expected, the resulting model precisely predicts outlet compositions, namely, with a relative error of 0.55% resulting from the prediction of 500 previously unseen inlet conditions in the training range. However, even positions within the reactor are predicted well, although no data were given for those positions. Figure 13a exemplarily shows the predicted composition profiles at a temperature of 655 K and a H_2/CO_2 ratio of 2.65. While overall accurate, the highest prediction error occurs for CO at a position around 4 mm (16.7%). This position is the maximum of the CO mass fraction. Consequently, the relative rate of change for CO mass fractions at that position is zero and falls below the threshold value of the physics loss. Such a tendency to under- or overestimate the extrema is, however, expected for any model that has to rely on minimizing the error in the apparent rates of change to predict composition profiles within a reactor. The same tendency can already be observed for our previous model based on fixed inlet compositions.

Overall, the AC-CatalyticReactor-PINN provides accurate composition profiles, yielding 4.5% relative error for compositions within the reactor and 2% relative error for extrapolation of up to 50% beyond the reactor length (Figure 13b,c).

Numerical Speed-Up

The computational speed-up achieved by the AC-CatalyticReactor-PINN relative to established reactor model solvers like Cantera depends strongly on the spatial resolution requested. For practically relevant resolutions of up to 100 axial positions, the speed-up exceeds 1000 (Figure 14). The

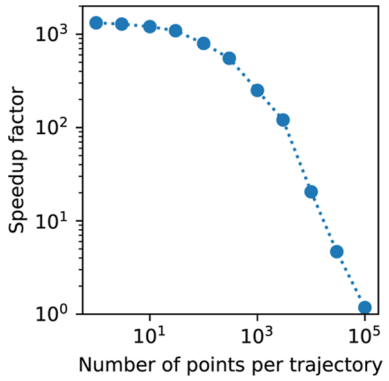

Figure 14. Numerical speed-up to obtain reactor compositions with the PINN surrogate compared to the classical solution of the reactor model using the Cantera software. The value depends on the spatial resolution requested along the reactor.

average solution time of the PFR model with a Cantera is 28 ms, while the AC-CatalyticReactor-PINN takes 21 μs in PyTorch per single point on a trajectory. Both values are measured on an AMD Ryzen 7 9800 X3D processor.

Several factors influence the exact performance gain, including the hardware specifications, the amount of utilized processor cores, the stiffness of the underlying ODE system, and the neural network size required to meet accuracy demands. In this study, two processor cores were used. The observed speed-up trends align with theoretical expectations: for large batch sizes, the scaling is approximately linear. For small batch sizes, however, the numerical overhead associated with PyTorch tensor initializations and memory overhead dominates the prediction time, resulting in a performance plateau.

While this work considers mass transport only through an idealized PFR model, additional effects, including axial and radial diffusion as well as heat transport, are easily implemented into the physics loss function. In these cases, an even higher speed-up can be expected because the PINN solution is meshless and does not require direct integration within spatial discretizations of the domain.

CONCLUSION

In this work, we introduced the AC-CatalyticReactor-PINN, a PINN architecture that enables accurate, meshless, and physically consistent modeling of catalytic reactors with detailed surface kinetics in real time.

Central to our approach is the enforcement of atom conservation through a dedicated neural network layer. We systematically evaluated three strategies to impose this constraint and demonstrated that the completion-based approach outperforms projection approaches in both accuracy and computational efficiency. Our proposed implementation ensures flexibility and ease of integration into any neural network-based application. In future work, this can also be extended to enforce the nonlinear energy conservation in nonisothermal systems by pairing the nonorthogonal projection steps explored in this work with iterative procedures like the ENFORCE method.[32]

We further addressed a major challenge in physics-informed modeling of catalytic systems: the incompatibility of detailed surface kinetic mechanisms with gradient-based PINN training. By employing a physically consistent surrogate model for steady-state source term evaluations, we re-enabled the use of automatic differentiation and accelerated training without compromising on physical fidelity. This surrogate is required only during training of AC-CatalyticReactor-PINN, preserving overall deployment simplicity.

Applying our model to a nickel-catalyzed methanation monolith reactor, we achieved accurate composition profiles using only 40 inlet−outlet data pairs and 4000 collocation points, demonstrating exceptional data efficiency. The trained model evaluates reactor trajectories 1000 times faster than established solvers such as Cantera, enabling use cases that demand real-time inference.

The AC-CatalyticReactor-PINN opens new opportunities in catalytic system design, embedded control, and parameter estimation from sparse measurements. Future work may extend this framework to dynamic reactor operation, more complex reactor models, or the integration into computational fluid dynamics solvers to accelerate transient reactive flow simulations. It lays the groundwork for a new class of catalytic reactor surrogates that are both physically grounded and computationally efficient, paving the way for real-time, high-fidelity simulations in chemical reaction engineering.

ASSOCIATED CONTENT

Supporting Information

The Supporting Information is available free of charge at https://pubs.acs.org/doi/10.1021/acsengineeringau.5c00106.

Reactor model description; microkinetic model of methanation; assessment of the accuracy of the kinetic surrogate; validity of the weighted projection method;

https://doi.org/10.1021/acsengineeringau.5c00106
ACS Eng. Au 2026, 6, 359−371

correction matrices for methanation reaction; choosing the key species; gradient propagation through steady-state surface kinetics; and comparison between the physics loss ramp and an adaptive weighting algorithm (PDF)

AUTHOR INFORMATION

Corresponding Authors

Mauro Bracconi − *Laboratory of Catalysis and Catalytic Processes, Dipartimento di Energia, Politecnico di Milano, Milano 20156, Italy;* orcid.org/0000-0001-7643-3214; Email: mauro.bracconi@polimi.it

Matteo Maestri − *Laboratory of Catalysis and Catalytic Processes, Dipartimento di Energia, Politecnico di Milano, Milano 20156, Italy;* orcid.org/0000-0002-8925-3869; Email: matteo.maestri@polimi.it

Author

Felix A. Döppel − *Laboratory of Catalysis and Catalytic Processes, Dipartimento di Energia, Politecnico di Milano, Milano 20156, Italy;* orcid.org/0000-0003-4733-9872

Complete contact information is available at: https://pubs.acs.org/10.1021/acsengineeringau.5c00106

Notes

The authors declare no competing financial interest.

ACKNOWLEDGMENTS

The authors thank the Seal of Excellence Programme by Politecnico di Milano, and F.D. thanks the Walter Benjamin Programme (grant no. 569475041) by the German Research Foundation (DFG) for funding this work.

ADDITIONAL NOTES

[1]To avoid numerical issues in the evaluation of eq 10, we suggest to replace zero weights on the diagonal by a small finite value of 10^{-12}.

[2]This means that instead of computing the correction matrix via eq 10, it can be obtained via calls to readily available functions in many programming languages.

REFERENCES

(1) Biermann, F.; Uglietti, R.; Döppel, F. A.; Kircher, T.; Maestri, M.; Votsmeier, M. *Chem. Eng. J.* **2025**, *519*, 163598.
(2) Karniadakis, G. E.; Kevrekidis, I. G.; Lu, L.; Perdikaris, P.; Wang, S.; Yang, L. *Nat. Rev. Phys.* **2021**, *3*, 422−440.
(3) Baydin, A. G.; Pearlmutter, B. A.; Radul, A. A.; Siskind, J. M. *J. Mach. Learn. Res.* **2018**, *18*, 1−43.
(4) Raissi, M.; Perdikaris, P.; Karniadakis, G. *J. Comput. Phys.* **2019**, *378*, 686−707.
(5) Ngo, S. I.; Lim, Y.-I. *IFAC-PapersOnLine* **2022**, *55*, 429−434.
(6) Lastrucci, G.; Theisen, M. F.; Schweidtmann, A. M. In *Comput.-Aided Chem. Eng.*; Elsevier Masson SAS, 2024; Vol. *53*, pp 571−576.
(7) Lastrucci, G.; Karia, T.; Gromotka, Z.; Schweidtmann, A. M. *arXiv* **2025**, arXiv:2501.17782.
(8) Spanò, G.; Ferri, M.; Cheula, R.; Monai, M.; Weckhuysen, B. M.; Maestri, M. *ACS Catal.* **2025**, *15*, 8194−8203.
(9) Gusmão, G. S.; Retnanto, A. P.; Cunha, S. C. d.; Medford, A. J. *Catal. Today* **2023**, *417*, 113701.
(10) Gusmão, G. S.; Medford, A. J. *Comput. Chem. Eng.* **2024**, *181*, 108547.
(11) Nai, D.; Gusmão, G. S.; Kilwein, Z. A.; Boukouvala, F.; Medford, A. J. *Digit. Discov.* **2024**, *3*, 2327−2340.
(12) Deutschmann, O. In *Handbook of Heterogeneous Catalysis*, Ertl, G., Knözinger, H., Schüth, F., Weitkamp, J., Eds., 2nd: ed.; Wiley VCH: 2008; Chapter Computational, pp 1811−1821.
(13) Votsmeier, M. *Chem. Eng. Sci.* **2009**, *64*, 1384−1389.
(14) Votsmeier, M.; Scheuer, A.; Drochner, A.; Vogel, H.; Gieshoff, J. *Catal. Today* **2010**, *151*, 271−277.
(15) Partopour, B.; Dixon, A. G. *Comput. Chem. Eng.* **2016**, *88*, 126−134.
(16) Partopour, B.; Dixon, A. G. *AIChE J.* **2017**, *63*, 87−94.
(17) Matera, S.; Maestri, M.; Cuoci, A.; Reuter, K. *ACS Catal.* **2014**, *4*, 4081−4092.
(18) Lorenzi, J. M.; Stecher, T.; Reuter, K.; Matera, S. *J. Chem. Phys.* **2017**, *147*, 164106.
(19) Sutton, J. E.; Lorenzi, J. M.; Krogel, J. T.; Xiong, Q.; Pannala, S.; Matera, S.; Savara, A. *ACS Catal.* **2018**, *8*, 5002−5016.
(20) Partopour, B.; Paffenroth, R. C.; Dixon, A. G. *Comput. Chem. Eng.* **2018**, *115*, 286−294.
(21) Bracconi, M.; Maestri, M. *Chem. Eng. J.* **2020**, *400*, 125469.
(22) Döppel, F.; Wenzel, T.; Herkert, R.; Haasdonk, B.; Votsmeier, M. *Chem. Ing. Tech.* **2024**, *96*, 759−768.
(23) Döppel, F. A.; Votsmeier, M. *Chem. Eng. Sci.* **2022**, *262*, 117964.
(24) Döppel, F. A.; Votsmeier, M. *React. Chem. Eng.* **2023**, *8*, 2620−2631.
(25) Klumpers, B.; Luijten, T.; Gerritse, S.; Hensen, E.; Filot, I. *Chem. Eng. J.* **2023**, *475*, 145538.
(26) Kircher, T.; Döppel, F. A.; Votsmeier, M. *Chem. Eng. J.* **2024**, *485*, 149863.
(27) Kircher, T.; Döppel, F. A.; Votsmeier, M. *Comput.-Aided Chem. Eng.* **2024**, 817−822.
(28) Fedorov, A.; Perechodjuk, A.; Linke, D. *Chem. Eng. J.* **2023**, *477*, 146869.
(29) Lacerda de Oliveira Campos, B.; Oliveira Souza da Costa, A.; Herrera Delgado, K.; Pitter, S.; Sauer, J.; Ferreira da Costa Junior, E. *React. Chem. Eng.* **2024**, *9*, 1047−1060.
(30) Kasiraju, S.; Vlachos, D. G. *React. Chem. Eng.* **2024**, *9*, 119−131.
(31) Döppel, F. A.; Votsmeier, M. *Proceedings of the Combustion Institute*; Elsevier, 2024; Vol. *40*, pp 105507.
(32) Lastrucci, G.; Schweidtmann, A. M. *arXiv* **2025**, arXiv:2502.06774.
(33) Min, Y.; Azizan, N. *arXiv* **2025**, arXiv:2410.10807.
(34) Chen, Y.; Huang, D.; Zhang, D.; Zeng, J.; Wang, N.; Zhang, H.; Yan, J. *J. Comput. Phys.* **2021**, *445*, 110624.
(35) Chen, H.; Flores, G. E.; Li, C. *Comput. Chem. Eng.* **2024**, *189*, 108764.
(36) Wang, T.; Yi, Y.; Yao, J.; Xu, Z.-q. J.; Zhang, T.; Chen, Z. *Combust. Flame* **2025**, *275*, 114105.
(37) Kircher, T.; Votsmeier, M. *J. Phys. Chem. Lett.* **2025**, *16*, 4715−4723.
(38) Beucler, T.; Pritchard, M.; Rasp, S.; Ott, J.; Baldi, P.; Gentine, P. *Phys. Rev. Lett.* **2021**, *126*, 98302.
(39) Liu, D. C.; Nocedal, J. *Math. Program.* **1989**, *45*, 503−528.
(40) Gao, B.; Yao, R.; Li, Y. *Comput. Math. Appl.* **2025**, *181*, 216−227.
(41) Schmider, D.; Maier, L.; Deutschmann, O. *Ind. Eng. Chem. Res.* **2021**, *60*, 5792−5805.
(42) Döppel, F.; Kircher, T.; Votsmeier, M. *ChemRxiv* **2025**, chemrxiv-2025-tz9fr.
(43) Goodwin, D. G.; Moffat, H. K.; Schoegl, I.; Speth, R. L.; Weber, B. W. Cantera: An object-oriented software toolkit for chemical kinetics, thermodynamics, and transport processes. *Zenodo* **2023**.
(44) Rohrhofer, F. M.; Posch, S.; Gößnitzer, C.; García-Oliver, J. M.; Geiger, B. C. *Energy and AI* **2024**, *16*, 100341.
(45) Sturm, P. O.; Silva, S. J. *ACS ES&T Air* **2025**, *2*, 99−108.