
Surrogate-Assisted PINNs with Hard Constraints for Heterogeneous Catalytic Reactor Modeling

Felix A. Döppel
Politecnico di Milano
via La Masa 34, 20156 Milano, Italy
felix.doeppel@polimi.it

Mauro Bracconi
Politecnico di Milano
via La Masa 34, 20156 Milano, Italy
mauro.bracconi@polimi.it

Matteo Maestri
Politecnico di Milano
via La Masa 34, 20156 Milano, Italy
matteo.maestri@polimi.it

Abstract

We propose a hard-constrained PINN framework for efficient catalytic reactor modeling that guarantees atom conservation through a dedicated neural network layer. By choosing the weights of this layer based on the concept of key species, we replace multiple output-nodes with physical constraints, while simultaneously preserving positivity and improving training stability. Further, we include a detailed micro kinetic description of the surface chemistry through a physically plausible Global Reaction Neural Network surrogate. Applied to a CO₂ methanation reactor, our approach achieves 1000× speed-up over conventional solvers while maintaining physical fidelity [1].

1 Introduction

Heterogeneous catalysts are directly involved in 85% of the large-scale chemical processes and play a major role for climate protection through the production of renewable energies and the removal of harmful pollutants [2]. Detailed numerical modeling lays the foundation for the rational design and on-line control of these devices [3]. In particular, multiscale modeling has shown great potential due to its ability to describe the physical phenomena across all relevant time and length scales [4]. However, this level of detail comes at the expense of high computational demand, often resulting in simulation times that are unbearable for industrial applications. Here, we propose hard constrained physics-informed neural networks (PINNs) [5, 6] as efficient and physically consistent surrogates of catalytic reactors with a detailed description of the surface chemistry [1].

In our atom conserving catalytic reactor PINN framework we, leverage the recently developed Global Reaction Neural Network (GRNN) [7] as a differentiable and physically plausible surrogate for detailed surface kinetics to re-enable stable gradient propagation via automatic differentiation. Otherwise, detailed surface kinetics would require solving a set of implicit functions which is not natively compatible with gradient computation via automatic differentiation. Further, we overcome the drawbacks of soft-constraining atom conservation, which can lead to residual violations that accumulate and destabilize simulations [8] as well as downstream analyses. To do so, we generalize recently developed linear equality

constraints [9, 10] into an atom conservation layer and identify the most effective choice of physically interpretable weights to enforce atom conservation.

2 Related work

While PINNs have been used for catalytic reactor modeling with simplified descriptions of the surface chemistry [11–13], there is a lack of solutions considering detailed surface chemistry through micro kinetic models. In contrast to the simplified kinetics, they require finding the steady state solution to systems of famously stiff ODEs. Their computational cost is so high, that it sparked a long tradition of kinetic surrogate modeling [7, 14–33]. The required stiff ODE solvers are not natively compatible with automatic differentiation, but even the usage of differentiable solver implementations [34] or implicit differentiation [35] would be computationally too demanding as the ODE system has to be solved for every collocation point in every PINN training epoch. To reduce the computational demand and simultaneously re-enable automatic differentiation, instead we employ the state-of-the-art physically plausible kinetic surrogate [7] to provide efficient and differentiable solutions to the detailed surface chemistry during the training of the AC-CatalyticReactor-PINN.

Recently, PINNs have been hard-constrain to follow the fundamental law of atom conservation through orthogonal projection [10] or projection in scaled space [8, 36]. We generalize the projection step by introducing physically inspired weights and embed it as a neural network layer.

3 Atom Conserving Catalytic Reactor-PINN framework

We present the AC-CatalyticReactor-PINN framework [1] to accelerate the solution of catalytic reactor models (Fig. 1). It uses the mass fractions ω_0 of each chemical species in the reactor inlet (feed) and the feed temperature T_0 to predict the composition ω at any other position z inside the reactor. The hidden layers use \tanh activation while the output layer uses exponential activation to guarantee positive outputs while effectively capturing the widely different orders of magnitude typically present in ω [11]. The subsequent AC-layer enforces atom conservation through a linear projection step mapping inconsistent outputs $\hat{\omega}$ onto consistent $\tilde{\omega}$ with the pre-computed weights $I - C$ and bias $\omega_0 C$ [10], where I is the unit matrix and C is computed from the elemental mass fractions F that are known *a priori*.¹

To estimate the training data ω over all relevant orders of magnitude, the data loss $\mathcal{L}_{\text{data}}$ is given as the mean squared asinh distance, a generalized version of the thresholded relative error [15, 27]. The threshold value $\text{rref}_{\text{data}}$ modulates between relevant $\omega > \text{rref}_{\text{data}}$ and irrelevant $\omega < \text{rref}_{\text{data}}$ values.

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

The physics loss $\mathcal{L}_{\text{physics}}$ applies the same distance measure to the normalized reaction rates \dot{s}/ω given by the detailed surface chemistry. To enable automatic differentiation through the detailed surface chemistry, we employ a physically plausible surrogate [7] GRNN $\approx \dot{s}(\omega, T)$. The derivative $\dot{\tilde{\omega}}$ of $\tilde{\omega}$ with respect to z is obtained through automatic differentiation.

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

¹ $C = (WW^T)^{-1} F^T (F(WW^T)^{-1} F^T)^{-1} F$ with weights $W = I$ for orthogonal projection, $W = \text{diag}(\hat{\omega}^{-1})$ for weighted projection, and $W_{i,i} = \begin{cases} 1 & \text{if } i \text{ is key species,} \\ 0 & \text{if } i \text{ is dependent species.} \end{cases}$ for completion.




Figure 1: Schematic representation of the 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.

The total loss $\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{data}} + \lambda \mathcal{L}_{\text{physics}}$ includes the weight λ to balance the typically unequal contributions of physics and data. While we used an empirical schedule for λ , this can easily be replaced by adaptive loss weighting algorithms [37].

4 Experiments

The AC-CatalyticReactor-PINN [1] is trained with a synthetic dataset consisting of inlet and outlet compositions of a 1D pseudo-homogeneous CO_2 methanation reactor model. The reactor is assumed isothermal and simulated with a molar inlet composition of 4% CO_2 , 5.3% H_2 , and 90.7% N_2 at 40 equally spaced temperatures in the range of 550 K to 850 K. The detailed surface kinetic scheme by SCHMIDER *et al.* [38] is employed and solved with IDAS [39]. 4000 collocation points are placed at uniformly random T and z values inside the training range.² Threshold values are chosen as $\text{rref}_{\text{data}} = 10^{-6}$ and $\text{rref}_{\text{physics}} = 5 \times 10^{-2} \text{ m}^{-1}$. PINN training is performed for 1000 epochs with the L-BFGS optimizer [40] using strong Wolfe line-search and a learning rate of 10^{-6} . Through preliminary experiments we found good convergence with $\lambda = 10^\alpha$ with α following a linear ramp of -8 to -2 from epoch 15 to 50.

5 Results and Discussion

We evaluate the AC-CatalyticReactor-PINN [1] performance by employing it as a surrogate for a methanation reactor model with detailed surface kinetics. The AC-CatalyticReactor-PINN accurately captures the spatial concentration profiles of all species when queried for unseen initial conditions within (Fig. 2a) or outside (Fig. 2b) the training range. Remarkably, it even captures the correct equilibrium composition (Fig. 2b). For a more general assessment, we consider 100 trajectories with unseen initial conditions and split them into two test sets. One with z values within the training region, and the other for up to 50% extrapolation. Here, the AC-CatalyticReactor-PINN achieves a relative prediction error of 3.8% for interpolation and 2.8% for extrapolation while atom balance is strictly enforced for any condition. These results are particularly promising given the compact training setup. The amount of training

²The rates obtained from every tenth IDAS solver step are saved and split into training/validation/test = 50/25/25 for training a GRNN with three hidden layers of 20 *tanh* nodes. Methane and hydrogen are assumed as key species. L-BFGS [40] with strong-wolfe line-search is used to reduce the thresholded relative error [15, 26] with $\text{rref} = 10^{-8} \text{ m}^{-1}$ until the validation error did not improve for 1000 consecutive epochs. The test error at unseen conditions is 0.25%.

data is experimentally well accessible and with fewer than 1000 parameters the model remains lightweight and suitable for embedded or resource-constrained deployment scenarios. For example, a computational speed-up of 1000 is reached compared to the established Cantera solver with a spatial resolution of 100 axial positions (21 μ s vs 28 ms on average).³

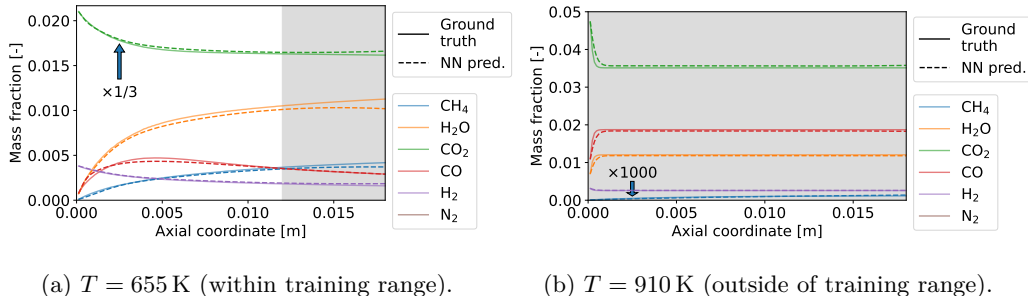


Figure 2: Composition trajectories predicted at unseen conditions (dashed) vs. ground truth (solid lines). Gray areas indicate extrapolation conditions.

The atom conservation layer is strongly contributing to the accuracy and robustness of the AC-CatalyticReactor-PINN. Omitting it leads to physically inconsistent results, with atom balance errors of up to 10 %, which is unacceptable for practical applications. Within the AC-layer, we explored different sets of weights, each corresponding to a different direction of the performed projection step. The orthogonal projection [10] maintains the atom balance but predicts negative concentrations for 1 % and 21.6 % for interpolation and extrapolation conditions, respectively. Such negative values are not only physically implausible but also incompatible with the kinetic surrogate, requiring cutoffs that undermine model convergence. The weighted projection aims to apply changes of equal relative magnitude for all chemical species. Thereby, it preserves positivity but its performance becomes depend on the prediction accuracy. In early training stages, inaccurate predictions lead to impractical projection directions, amplifying deviations from the ground truth. Even a pre-training phase does not fully resolve the convergence issue. In contrast, completion [41] is less sensitive to inaccurate predictions and therefore improves convergence from the first epoch. Further, it reduces the number of output nodes by the amount of chemical elements in the system [42, 43] (here, it’s four) from six to two. This reduction in unphysical degrees of freedom prevents overfitting and ensures physically plausible inputs to the kinetic surrogate during PINN training.

During the PINN training, the kinetic surrogate (GRNN) is queried $>10^8$ times. Calling the full surface kinetics instead, training would take ~ 37 days. With the GRNN, training completes in 191 s, corresponding to a speed-up of $\sim 16\,500$.³ This acceleration is crucial: it turns previously infeasible training into a practical workflow on consumer-grade hardware.

6 Conclusions and Future Work

We proposed the AC-CatalyticReactor-PINN framework for efficient catalytic reactor modeling with full physical fidelity [1]. In particular, we were able to employ detailed surface kinetics through surrogate-assisted training. Not only did this re-enable gradient flow through automatic differentiation, but also vastly reduced the computational demand of the physics loss, accelerating model training by four orders of magnitude. Comparing different atom conservation hard-constraint implementations, we identified the completion strategy as the most promising. With this, positivity is preserved and unphysical degrees of freedom are removed. We found stable training and a 4-fold increase in accuracy.

Using an isothermal CO₂ methanation showcase with only 40 training data, we achieved a 1000 \times speed-up with relative errors below 4%. Importantly, the framework is readily extendable to non-isothermal systems by including the energy balance in the physics loss.

³Data collected using 32 GB of RAM @ 6000 MHz and a Ryzen 7 9800X3D processor @ 5.2 GHz.

Overall, the AC-CatalyticReactor-PINN paves the way for efficient yet accurate industrial-scale catalytic reactor modeling, enabling real-time control and on-line process optimization. Moreover, we anticipate that the surrogate-assisted training strategy could also benefit other domains considering multi-scale systems involving complex surface kinetics, including atmospheric chemistry [44, 45], battery research [46–48], and more [49, 50].

Acknowledgments and Disclosure of Funding

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.

References

- [1] F. A. Döppel, M. Bracconi, M. Maestri, *ACS Engineering Au* **2026**, DOI 10.1021/acseengineeringau.5c00106.
- [2] W. Reschetilowski, *Einführung in die Heterogene Katalyse*, Springer-Verlag, Berlin, Heidelberg, **2015**.
- [3] D. Agar, M. Bertau, M. Börnhorst, M. Busch, M. Casapu, P. Claus, K. HerreraDelgado, D. Demtröder, O. Deutschmann, R. Dittmeyer, C. Dreiser, D. Eckes, K. Ehrhardt, B. Etzold, G. Fieg, H. Freund, J. Friedland, J.-D. Grunwaldt, R. Güttel, E. vonHarbou, J. Heck, T. Herrmann, K.-O. Hinrichsen, R. Horn, J. Khinast, E. Klemm, N. Kockmann, G. Kolios, U. Krewer, R. Kuwertz, M. Löffelholz, G. Luinstra, M. Nilles, H. Marschall, L. Möltner, M. Muhler, U. Nieken, M. Nilles, J. Osiewacz, S. Palkovits, M. Paul, K. Pflug, S. Pinnow, S. Radl, J. Sauer, V. Schallhart, S. Schunk, A. Seidel-Morgenstern, M. Schlüter, M. Schubert, K. Sundmacher, T. Turek, I. Vittorias, H. Vogel, O. Wachsen, G. Wehinger, H.-W. Zanthoff, D. Ziegenbalg, *Roadmap Chemical Reaction Engineering, Vol. 3*, DECHEMA/VDI Subject Division Chemical Reaction Engineering, **2023**.
- [4] G. D. Wehinger, M. Ambrosetti, R. Cheula, Z.-B. Ding, M. Isoz, B. Kreitz, K. Kuhlmann, M. Kutscherauer, K. Niyogi, J. Poissonnier, R. Réocreux, D. Rudolf, J. Wagner, R. Zimmermann, M. Bracconi, H. Freund, U. Krewer, M. Maestri, *Chemical Engineering Research and Design* **2022**, *184*, 39–58.
- [5] M. Raissi, P. Perdikaris, G. Karniadakis, *Journal of Computational Physics* **2019**, *378*, 686–707.
- [6] G. E. Karniadakis, I. G. Kevrekidis, L. Lu, P. Perdikaris, S. Wang, L. Yang, *Nature Reviews Physics* **2021**, *3*, 422–440.
- [7] T. Kircher, F. A. Döppel, M. Votsmeier, *Chemical Engineering Journal* **2024**, *485*, 149863.
- [8] T. Wang, Y. Yi, J. Yao, Z.-q. J. Xu, T. Zhang, Z. Chen, *Combustion and Flame* **2025**, *275*, 114105.
- [9] Y. Chen, D. Huang, D. Zhang, J. Zeng, N. Wang, H. Zhang, J. Yan, *Journal of Computational Physics* **2021**, *445*, 110624.
- [10] H. Chen, G. E. Flores, C. Li, *Computers and Chemical Engineering* **2024**, *189*, 108764.
- [11] S. I. Ngo, Y.-I. Lim, *IFAC-PapersOnLine* **2022**, *55*, 429–434.
- [12] G. Lastrucci, M. F. Theisen, A. M. Schweidtmann in *Computer Aided Chemical Engineering, Vol. 53*, Elsevier Masson SAS, **2024**, pp. 571–576.
- [13] G. Lastrucci, T. Karia, Z. Gromotka, A. M. Schweidtmann, *arXiv preprint arXiv:2501.17782* **2025**.
- [14] M. Votsmeier, *Chemical Engineering Science* **2009**, *64*, 1384–1389.
- [15] M. Votsmeier, A. Scheuer, A. Drochner, H. Vogel, J. Gieshoff, *Catalysis Today* **2010**, *151*, 271–277.
- [16] B. Partopour, A. G. Dixon, *Computers & Chemical Engineering* **2016**, *88*, 126–134.
- [17] B. Partopour, A. G. Dixon, *AIChE Journal* **2017**, *63*, 87–94.
- [18] M. Klingenberger, O. Hirsch, M. Votsmeier, *Computers & Chemical Engineering* **2017**, *98*, 21–30.

- [19] S. Matera, M. Maestri, A. Cuoci, K. Reuter, *ACS Catalysis* **2014**, *4*, 4081–4092.
- [20] J. M. Lorenzi, T. Stecher, K. Reuter, S. Matera, *Journal of Chemical Physics* **2017**, *147*, 164106.
- [21] J. E. Sutton, J. M. Lorenzi, J. T. Krogel, Q. Xiong, S. Pannala, S. Matera, A. Savara, *ACS Catalysis* **2018**, *8*, 5002–5016.
- [22] R. Uglietti, M. Bracconi, M. Maestri, *Reaction Chemistry & Engineering* **2020**, *5*, 278–288.
- [23] B. Partopour, R. C. Paffenroth, A. G. Dixon, *Computers & Chemical Engineering* **2018**, *115*, 286–294.
- [24] M. Bracconi, M. Maestri, *Chemical Engineering Journal* **2020**, *400*, 125469.
- [25] F. Döppel, T. Wenzel, R. Herkert, B. Haasdonk, M. Votsmeier, *Chemie Ingenieur Technik* **2024**, *96*, 759–768.
- [26] F. A. Döppel, M. Votsmeier, *Chemical Engineering Science* **2022**, *262*, 117964.
- [27] F. A. Döppel, M. Votsmeier, *Reaction Chemistry & Engineering* **2023**, *8*, 2620–2631.
- [28] B. Klumpers, T. Luijten, S. Gerritse, E. Hensen, I. Filot, *Chemical Engineering Journal* **2023**, *475*, 145538.
- [29] A. Fedorov, A. Perechodjuk, D. Linke, *Chemical Engineering Journal* **2023**, *477*, 146869.
- [30] S. Kasiraju, D. G. Vlachos, *Reaction Chemistry & Engineering* **2024**, *9*, 119–131.
- [31] T. Kircher, F. A. Döppel, M. Votsmeier in *Computer Aided Chemical Engineering*, **2024**, pp. 817–822.
- [32] B. Lacerda de Oliveira Campos, A. Oliveira Souza da Costa, K. Herrera Delgado, S. Pitter, J. Sauer, E. Ferreira da Costa Junior, *Reaction Chemistry & Engineering* **2024**, *9*, 1047–1060.
- [33] F. Biermann, R. Uglietti, F. A. Döppel, T. Kircher, M. Bracconi, M. Maestri, M. Votsmeier, *Chemical Engineering Journal* **2025**, 163598.
- [34] R. T. Q. Chen, Y. Rubanova, J. Bettencourt, D. K. Duvenaud, *Advances in Neural Information Processing Systems* **2018**, *31*.
- [35] S. Bai, J. Z. Kolter, V. Koltun, *Advances in Neural Information Processing Systems* **2019**, *32*.
- [36] T. Kircher, M. Votsmeier, *The Journal of Physical Chemistry Letters* **2025**, *16*, 4715–4723.
- [37] B. Gao, R. Yao, Y. Li, *Computers Mathematics with Applications* **2025**, *181*, 216–227.
- [38] D. Schmider, L. Maier, O. Deutschmann, *Industrial and Engineering Chemistry Research* **2021**, *60*, 5792–5805.
- [39] A. C. Hindmarsh, P. N. Brown, K. E. Grant, S. L. Lee, R. Serban, D. E. Shumaker, C. S. Woodward, *ACM Transactions on Mathematical Software* **2005**, *31*, 363–396.
- [40] D. C. Liu, J. Nocedal, *Mathematical Programming* **1989**, *45*, 503–528.
- [41] T. Beucler, M. Pritchard, S. Rasp, J. Ott, P. Baldi, P. Gentine, *Physical Review Letters* **2021**, *126*, 98302.
- [42] M. Baerns, A. Behr, H. Hofmann, J. Gmehling, U. Onken, A. Renken, K.-O. Hinrichsen, R. Palkovits, *Technische Chemie*, Wiley-VCH, Weinheim, **2013**.
- [43] F. A. Döppel, M. Votsmeier, *Proceedings of the Combustion Institute* **2024**, *40*, 105507.
- [44] G. C. Sosso, T. Li, D. Donadio, G. A. Tribello, A. Michaelides, *The Journal of Physical Chemistry Letters* **2016**, *7*, 2350–2355.
- [45] R. Signorell, A. Bertram, *Physical Chemistry Chemical Physics* **2009**, *11*, 7759.
- [46] B. Liu, C. Wei, Z. Zhu, Y. Fang, Z. Bian, X. Lei, Y. Zhou, C. Tang, Y. Qian, G. Wang, *Angewandte Chemie International Edition* **2022**, *61*, e202212780.
- [47] J. Zheng, Y. Deng, W. Li, J. Yin, P. J. West, T. Tang, X. Tong, D. C. Bock, S. Jin, Q. Zhao, R. Garcia-Mendez, K. J. Takeuchi, E. S. Takeuchi, A. C. Marschilok, L. A. Archer, *Science Advances* **2022**, *8*.
- [48] J. Li, X. Xiao, Y.-T. Cheng, M. W. Verbrugge, *The Journal of Physical Chemistry Letters* **2013**, *4*, 3387–3391.

- [49] H. M. Cuppen, C. Walsh, T. Lamberts, D. Semenov, R. T. Garrod, E. M. Penteado, S. Ioppolo, *Space Science Reviews* **2017**, *212*, 1–58.
- [50] O. Deutschmann, R. Schmidt, F. Behrendt, J. Warnat, *Symposium (International) on Combustion* **1996**, *26*, 1747–1754.