This is an electronic reprint of the original article.
This reprint may differ from the original in pagination and typographic detail.

Haitsiukevich, Katsiaryna; Bergman, Samuli; de Araujo Filho, Cesar ; Corona, Francesco; Ilin, Alexander

A Grid-Structured Model of Tubular Reactors

Published in:
2021 IEEE 19th International Conference on Industrial Informatics (INDIN)

DOI:
[10.1109/INDIN45523.2021.9557382](https://doi.org/10.1109/INDIN45523.2021.9557382)

Published: 21/07/2021

Document Version
Peer reviewed version

Please cite the original version:
Haitsiukevich, K., Bergman, S., de Araujo Filho, C., Corona, F., & Ilin, A. (2021). A Grid-Structured Model of Tubular Reactors. In *2021 IEEE 19th International Conference on Industrial Informatics (INDIN)* IEEE. <https://doi.org/10.1109/INDIN45523.2021.9557382>

This material is protected by copyright and other intellectual property rights, and duplication or sale of all or part of any of the repository collections is not permitted, except that material may be duplicated by you for your research use or educational purposes in electronic or print form. You must obtain permission for any other use. Electronic or print copies may not be offered, whether for sale or otherwise to anyone who is not an authorised user.

A Grid-Structured Model of Tubular Reactors

Katsiaryna Haitisiukevich Samuli Bergman, Cesar de Araujo Filho Francesco Corona, Alexander Ilin
Aalto University Neste, Espoo, Finland Aalto University
Espoo, Finland {firstname.lastname}@neste.com Espoo, Finland
{firstname.lastname}@aalto.fi {cesar.araujo}@neste.com {firstname.lastname}@aalto.fi

Abstract—We propose a grid-like computational model of tubular reactors. The architecture is inspired by the computations performed by solvers of partial differential equations which describe the dynamics of the chemical process inside a tubular reactor. The proposed model may be entirely based on the known form of the partial differential equations or it may contain generic machine learning components such as multi-layer perceptrons. We show that the proposed model can be trained using limited amounts of data to describe the state of a fixed-bed catalytic reactor. The trained model can reconstruct unmeasured states such as the catalyst activity using the measurements of inlet concentrations and temperatures along the reactor.

Index Terms—catalyst activity, deep learning, fixed-bed catalytic reactor, multi-layer perceptron, soft sensor, tubular reactor

I. INTRODUCTION

Tubular reactors are commonly used in process industry for producing value-added products by operating a chemical reaction. During operation, the reactor is constantly monitored to control the process inside the reactor, to detect malfunctioning or to schedule maintenance. It is typical that many variables describing the state of the reactor cannot be directly measured during operation and they must be estimated using measured variables. This is a classical state estimation problem which requires a dynamic model of the reactor.

There is extensive literature on modeling tubular reactors using first principles, e.g. [1], [2]. The chemical process inside the reactor is usually described by partial differential equations (PDEs). The parameters of the PDEs are determined by conservation laws and in the ideal world they would fit the observed data well. Unfortunately, due to inaccuracies in modeling assumptions and high level of noise, the first-principles models often do not agree with the data. This motivates building reactor models in a data-driven way by fitting generic parametric models such as neural networks [3] to process observations. This is a viable solution in case one can collect large amounts of observations, otherwise the problem of overfitting becomes inevitable. Unfortunately, data scarcity is a common problem in this domain.

In this work, we assume that the tubular reactor system can be described by PDEs whose form is known but which contain some unknown nonlinearities. We propose a model whose computational graph is based on the knowledge of how PDEs are discretized and solved. The proposed model is a deep model with a grid structure in which each cell performs computations inspired by the known form of the PDEs. Each

cell may optionally contain generic neural networks to model unknown nonlinear components that can be present in the PDEs. The model parameters are estimated using gradient descent and backpropagation, which is common for deep learning models. We use simulated data to show that our model can learn to recover the state of a chemical process in a fixed bed reactor with catalyst poisoning. We apply the same model to data collected from a real chemical reactor and show that the model provides a good fit.

II. FIXED BED REACTOR WITH CATALYST POISONING

As an example tubular reactor, we consider a fixed bed reactor with exothermic reaction $A + H \rightarrow B$, in which a set of aromatic compounds collectively denoted by A are hydrogenated on the surface of the catalyst to produce a set of de-aromatized compounds B [4]. We assume that the catalyst deactivates irreversibly due to adsorption of poisoning agent P, meaning that the catalyst activity is decreasing over time. We also assume that there is a constant abundance of H.

A macroscopic characterization of this type of reactors in space and time t is given by a set of PDEs accounting for mass and energy balances. We restrict ourselves to a single spatial dimension, the axis of the reactor z . We also assume that the reactor operates under plug flow conditions, that is the diffusion phenomena are negligible, and consider only convection and reaction phenomena.

The concentration $x_a(z, t)$ of aromatic compounds A in the reactor is given as a solution to the mass balance equation

$$\frac{\partial x_a}{\partial t} = -U \frac{\partial x_a}{\partial z} - \alpha(T)r_a, \quad (1)$$

where U is the constant velocity of the fluid along the reactor and r_a is the reaction rate at which A is converted into B. The term $\alpha(T)$ denotes a linear function of temperature T .

Similarly, the concentration $x_p(z, t)$ of poison P solves

$$\frac{\partial x_p}{\partial t} = -U \frac{\partial x_p}{\partial z} - \alpha(T)r_p, \quad (2)$$

where r_p is the rate at which P is adsorbed on the catalyst.

Under the assumption that the reactor is adiabatic, values $T(z, t)$ of the temperature function solve the energy balance

$$\frac{\partial T}{\partial t} = -\beta(x_a, T)U \frac{\partial T}{\partial z} + \gamma r_a, \quad (3)$$

where $\beta(x_a, T)$ is proportional to x_a/T and γ is a constant.

Fig. 1. One catalyst cycle of the simulated data. Each color represents one reactor level i . The signals are scaled to the range $[0, 1]$.

To account for decay in catalyst activity, $\theta(z, t)$ solves

$$\frac{d\theta}{dt} = -r_d, \quad (4)$$

where r_d is the deactivation rate. For all locations, the catalyst activity starts at one and decays to zero with time.

The rate r_a of the hydrogenation reaction is

$$r_a = \frac{k_0 K_0 \exp\left(-\frac{Q-E}{RT}\right)}{1 + K_0 \exp\left(-\frac{Q}{RT}\right) P x_a^{l_2}} P^2 x_a^{l_1} x_h \theta, \quad (5)$$

where Q and E are adsorption and hydrogenation activation energies, k_0 is pre-exponential factor for hydrogenation, K_0 is aromatics adsorption constant, l_1 and l_2 are exponents constants, P is pressure, R is gas constant and the concentration of hydrogen x_h which is assumed to be constant. The rate r_p at the which the poisoning agent P is adsorbed onto the catalyst and the catalyst deactivation rate r_d are related:

$$r_d = k_d^0 \exp(-E_d/(RT)) P x_p \theta \quad (6)$$

$$r_p = r_d C_d \theta, \quad (7)$$

where E_d is activation energy for the catalyst poisoning, k_d^0 is poisoning pre-exponential factor and C_d is the catalyst poison adsorption capacity.

For a given set of initial and boundary conditions, Fig. 1 shows the time and space evolution of the system variables over one catalyst cycle. At the beginning of the cycle, the reaction and catalyst deactivation occur close to the reactor inlet, as indicated by the large spatial gradients of x_a , T and x_p . Later, catalyst poisoning causes the shift of the reaction towards the reactor outlet. When the catalyst is deactivated, the concentrations x_a , x_p and the temperatures T do not change any longer and the reactor behaves like a tubular pipe.

III. GRID-STRUCTURED MODEL FOR TUBULAR REACTORS

To design a surrogate neural model of the reactor, we assume that the process develops slowly in time, that is the reactor is in a pseudo-steady-state. We set the time derivatives

Fig. 2. Left: A plan of a tubular reactor with the locations of the temperature sensors. Black circles indicate the location without sensors. Right: The computational graph of the proposed model where $\mathbf{x}$ denotes a vector with components x_a, x_p, T .

$\frac{\partial x_a}{\partial t}$, $\frac{\partial x_p}{\partial t}$, $\frac{\partial T}{\partial t}$ to zero which modifies (1)–(4) to a set of ordinary differential equations (ODEs). Assuming that we know the concentrations $x_a(z_0, t)$, $x_p(z_0, t)$ and temperatures $T(z_0, t)$ at the reactor inlet and that initial catalyst activity $\theta(z, t_0) = 1$, we can predict the concentrations, temperatures and the catalyst activity along the reactor by solving the resulting set of ODEs using the Euler's method:

$$x(z + \Delta z, t) \approx x(z, t) + f_x(x_a, x_p, T, \theta) \Delta z \quad (8)$$

$$\theta(z, t + \Delta t) \approx \theta(z, t) - r_d \Delta t, \quad (9)$$

where $x(z, t)$ denotes one of the three variables x_a , x_p or T and functions f_x are derived from (1)–(3) with partial derivatives $\frac{\partial x_a}{\partial t}$, $\frac{\partial x_p}{\partial t}$, $\frac{\partial T}{\partial t}$ set to zero. These computations can be implemented by a grid-structured computational graph shown in Fig. 2. Each cell in that graph corresponds to one point (z, t) in a chosen discretization of space and time. The inputs to the model are the boundary and initial conditions. The vertical top-down connections implement computations in (8) while the horizontal connections represent (9). The parameters of each cell are defined by the parameters of the original PDEs (1)–(4). The parameters are shared for all cells to guarantee the required property of equivariance to translations in space and time.

Suppose now that the model is unknown and we need to fit it to process observations. We assume that we can measure the concentrations $x_a(z_0, t)$, $x_p(z_0, t)$ at the inlet and the temperatures $T^*(z_i, t_j)$ at locations with indices $i \in I_z$ along the reactor at times with indices $j \in I_t$. We can tune the model by building a computational graph in Fig. 2 and by fitting the model outputs to the temperature observations. This can be done by minimizing the mean-squared error loss

$$C = \frac{1}{|I_z||I_t|} \sum_{i \in I_z} \sum_{t \in I_t} (T(z_i, t_j) - T^*(z_i, t_j))^2. \quad (10)$$

The proposed architecture can be seen as a deep recurrent model in which each cell corresponds one location in space

and time. Tuning the parameters of the model can be challenging due to vanishing and exploding gradients, like in other deep learning models. Standard recurrent units such as gated recurrent unit (GRU) [5] or long-short term memory (LSTM) [6], which are designed to mitigate the vanishing/exploding gradients problem, can be used as computational units in the proposed model. However, this comes with the price of losing the model interpretability and requiring much richer data to cover all possible operating conditions of the reactor.

IV. RELATED WORK

Several neural networks have architectures resembling or motivated by numerical solvers of differential equations (see, e.g., [7]–[10]). Many works built data-driven models by combining generic neural networks with PDE and ODE solvers (see, e.g., [10], [11]). Our model can be viewed as a particular instance of interactions networks [12] in which each level of the reactor interacts with the neighboring levels. Instead of using generic function approximators to model the interactions, we limit the family of functions to achieve interpretability and sample efficiency.

In the chemical engineering literature, several works proposed grid-structured architectures for modeling reactors [13], [14]. The model proposed in [13] was used for predicting the mixing characteristics of fixed beds of spheres in fixed-bed catalytic reactors. A similar approach was used in [14] to model the behavior of trickle-bed reactors employed for the exothermic hydrotreating of benzene. They described the steady-state of the reactor using one-dimensional or two-dimensional networks in which dimensions corresponded to the geometrical dimensions of the reactor. Three model parameters in [14] were learned from the data while the rest of the parameters were taken from literature. In our approach, the dimensions of the grid correspond to space and time. We learn all the model parameters in a gradient-based optimization.

V. EXPERIMENTS WITH SIMULATED DATA

A. Synthetic data

We first conducted experiments on synthetic data generated by solving PDEs (1)–(4) numerically. To simulate the data, we used constant temperature at the inlet, while the inlet concentrations were varied (see Fig. 1). The generated data were downsampled such that one trajectory representing one catalyst cycle contained 93 time points at 46 reactor levels. Fig. 1 shows the simulated data with one catalyst cycle that were used for training.

The inputs of the model were the concentrations and the temperature at the reactor inlet and temperature measurements $T(z_j, t)$ at locations $0, 4, 8, \dots, 44$ were used as the targets. Temperatures at locations $1, 5, \dots, 41$ were used as validation set 1. Temperatures at the remaining locations were used as validation set 2 that we used for model selection. Note that the validation sets significantly correlate with the training data because all data come from the same catalyst cycle. We nevertheless used such validation sets because in practice the modeler often has measurements of only one catalyst

cycle. We tested the model performance on five catalyst cycles simulated with the following scenarios. In Test 1, we used the same operating conditions as in the training data but different random seed. In Test 2, the inlet temperature was increased by 7% compared to the training scenario. In Test 3, the inlet concentrations x_a were increased by 20%. In Test 4, the inlet concentrations x_p were decreased by 15%. Test 5 combines the changes from Tests 2–4.

B. Training details

We trained the proposed grid-structured model in which each cell implemented the following computations:

$$x_a^{\text{out}} = x_a - \frac{k_a}{U} g_a(T, x_a) \theta \quad (11)$$

$$x_p^{\text{out}} = x_p - \frac{k_p}{U} T g_p(T, x_p) \theta^2 \quad (12)$$

$$T^{\text{out}} = T + \frac{k_T g_a(T, x_a)}{U(k_1 + k_2 x_a)} \theta \quad (13)$$

$$\theta^{\text{out}} = \theta - k_\theta g_p(T, x_p) \theta, \quad (14)$$

where x_a, x_p, T, θ are the cell inputs, $x_a^{\text{out}}, x_p^{\text{out}}, T^{\text{out}}, \theta^{\text{out}}$ are the cell outputs and $k_a, k_p, k_T, k_\theta, k_1, k_2$ are the model parameters. The functional form of the cell was derived from the PDEs (1)–(4). The important property of this cell is that the same functions g_a, g_p are used in (11), (13) and (12), (14), respectively. This couples the dynamics of the four variables and makes it possible to reconstruct the unobserved variables using the observed ones.

In the first version of the cell, the functional forms of g_a and g_p was dictated by the original PDE model $g_a(T, x_a) = e^{-k_3/T} T x_a$, $g_p(T, x_p) = e^{-k_4/T} x_p$. We refer to this version of the cell as *PDE param*. In the second version of the cell, we assumed that the exact forms of the reaction and adsorption rates were not known and therefore functions g_a, g_p were modeled with generic MLP networks. The MLPs had the architectures 32-64-32-1, 32-32-1 for g_a and g_p respectively, with layer normalization [15] followed by ReLU non-linearity after each hidden layer and softplus nonlinearity in the output layer. The MLP inputs were normalized with min-max normalization. We denote this cell *MLP* in the tables with the results.

When training the model with the MLP-based cell, we optionally used two extra penalty terms in the loss function:

$$C_1 = \lambda_a \sum_{i,j} \min(0, x_a(z_i, t_j))^2 + \lambda_p \sum_{i,j} \min(0, x_p(z_i, t_j))^2$$

$$C_2 = \lambda_a \sum_{t_j < t'} x_a(z_N, t_j)^2 + \lambda_p \sum_{t_j < t'} x_p(z_N, t_j)^2.$$

The first penalty term forces the non-negativity of the concentrations while the second term forces the outlet concentrations at the beginning of the catalyst cycle to be close to zero. We used hyperparameters $\lambda_a = 1$ and $\lambda_p = 100$ and $t' = 46$. We denote the model trained with such regularization as *MLP-reg*.

All MLP-based models were trained with the Adam optimizer [16] for 1000 gradient updates with learning rate 10^{-3} and then tuned for 200 more updates with learning rate 10^{-4} .

TABLE I
NORMALIZED RMSE OF TEMPERATURE RECONSTRUCTIONS IN EXPERIMENTS WITH SIMULATED DATA.

	T , training cycle			ΔT , training cycle			ΔT , test cycles				
	Train	Valid. set 1	Valid. set 2	Train	Valid. set 1	Valid. set 2	Test 1	Test 2	Test 3	Test 4	Test 5
GRU with 8 neurons	0.04	—	—	0.13	—	—	0.57	1.34	1.04	2.26	2.12
Grid-GRU	0.10	2.62	0.12	0.19	5.86	0.67	0.28	0.83	0.78	0.74	1.58
Grid-GRU with augm	0.14	0.14	0.14	0.25	0.35	0.30	0.27	0.99	0.36	0.43	1.26
PDE param	0.15	0.14	0.14	0.23	0.23	0.24	0.25	0.35	0.27	0.26	0.39
MLP	0.14	0.13	0.13	0.16	0.15	0.16	0.19	0.54	0.36	0.33	0.54
MLP-reg	0.13	0.12	0.12	0.17	0.17	0.17	0.21	0.38	0.26	0.27	0.38

The model with the PDE parametrization was trained with Adam with the same schedule but with ten times larger rates.

As the first baseline model, we used a generic GRU with eight neurons. The inputs of the GRU model were the inlet concentrations, the inlet temperature and velocity U . Naturally, this model can predict temperatures only at the same locations which were used in the training set. The second baseline was a grid model with the same computational graph as in Fig. 2, in which each cell was implemented by a generic GRU with eight neurons. Compared to the simple GRU baseline, this model has an in-built property of spatial equivariance. The inputs of both GRU-based models were normalized. Both baseline models were trained for 3000 epochs with learning rate 10^{-3} and then for 3000 epochs more with learning rate 10^{-4} .

C. Results

Table I shows the root mean-squared errors normalized by the data standard deviation (NRMSE) for the temperature reconstructions obtained with the trained models. We present the quality of the reconstructions of temperatures T and temperature differences between two subsequent sensor locations in a given set:

$$\Delta T(z_i, t) = T(z_i, t) - T(z_{j_i}, t), \quad (15)$$

where z_i and z_{j_i} are the locations of successive sensors. For example, for the training set we compute temperature differences between locations 4 and 0, 8 and 4 and so on. ΔT is an indicative metric because it reflects the reaction rate at a particular location. The best model for each class was selected from five runs with different seeds based on the MSE of ΔT reconstruction on validation set 2.

The GRU baseline model severely overfits to the training data and completely fails to generalize to the test cycles due to the lack of the right inductive bias. The grid-GRU model trained in a standard way finds a solution that overfits to the regular spacing of the temperature sensors. This is indicated by the large error on validation set 1. The networks learns to produce meaningful reconstructions only at the locations of the measurements, thus using intermediate locations of the grid as hidden layers that can develop arbitrary representations.

To mitigate this problem, we trained the grid-GRU model with the following data augmentation scheme. At each training iteration, we randomized the locations selected from the model prediction vector for the loss calculation by (10). Each location was shifted independently either to the previous level (with

probability 0.25) or to the next level (with probability 0.25) or kept at the correct location. The model trained in this way is denoted as *Grid-GRU with augm*. The results show that the Grid-GRU model trained with augmentations was able to generalize much better. The model performed well on Tests 1, 3 and 4 but failed on Tests 2 and 5, that is the model could not generalize to temperature changes.

The proposed grid-structured models showed very good performance on all of the five test sets. The best performance was obtained with the PDE-based cell. This model used the correct analytical form of the reaction and poisoning rates, which let it extrapolate well outside the training distribution. The model with a generic MLP cell performed very well too and it was able to outperform the PDE-based model on some tests.

TABLE II
AVERAGE $\bar{R}^2$ SCORE FOR x_a , x_p , T AND θ RECONSTRUCTIONS

Model	Train	Test 1	Test 2	Test 3	Test 4	Test 5
PDE param	0.963	0.958	0.959	0.959	0.957	0.956
MLP	0.953	0.943	0.945	0.955	0.895	0.913
MLP-reg	0.967	0.957	0.954	0.960	0.942	0.945

In Table II, we compare the accuracy of the grid-structured models in the task of predicting the values of the unobserved variables in all spatial and temporal locations. As the metric, we use the average of R^2 scores for four system states. Note that the table contains results only for the proposed grid models because the baseline models cannot reconstruct unobserved system states. The high values of the R^2 -score indicate that both models produced excellent reconstructions. The PDE-based model is generally slightly better than the MLP-based model. Fig. 3 shows the reconstructions of the catalyst activity and the temperature differences ΔT , as defined in (15), at different reactor levels. One can see that the produced estimates are close to the ground truth.

VI. EXPERIMENTS WITH REAL DATA

We applied the proposed models to estimate the states of a large-scale hydrotreatment unit. The unit is used for continuous saturation of hydrocarbons using hydrogen. The reductive hydrogenation occurs on the surface of a single bed of a catalyst. The reaction releases heat and the process is assumed to be adiabatic. The presence of a poisoning component and other impurities leads to irreversible catalyst deactivation.

Fig. 3. Experiments with simulated data: The catalyst activity (above) and the temperature differences ΔT (below) reconstructed at different reactor levels by the proposed grid-structured models. Left: Cell with PDE parameterization, middle: MLP-based cell, right: MLP-based cell with regularization. The dotted lines represent the simulated ground truth. The ΔT time series are min-max scaled.

Thus, the catalyst needs to be periodically replaced. The considered unit is equipped with two sequential reactors of the same nature, the outlet of the first reactor is directly fed to the second one. The unit is monitored by taking the following measurements:

- the flow rate and the temperature of the hydrocarbon reactor feed, measured in real time;
- the feed composition including the concentration of reactants and the poisoning agent, measured periodically by laboratory analysis;
- the temperatures inside the reactors, measured in real time by an array of thermocouples (sensors) installed inside the catalyst bed.

We used data from a period which represented one life cycle of the catalyst. Since our focus was on modeling slowly changing phenomena such as catalyst deactivation, we took daily averages of temperature measurements and downsampled them to create time series $T(z_i, t)$ with the same sampling frequency as the laboratory measurements of the concentrations $x_a(z_0, t)$ and $x_p(z_0, t)$. Linear velocity $U(t)$ was approximated by the flow rate of the hydrocarbon feed.

We built a model of the whole unit which included both reactors. We discretized the dimension along the flow direction such that the discretization points were close to the locations of the temperature sensors. This resulted in 46 locations z_i , where $i = 0$ and $i = 21$ were the locations of the inlets of the first and the second reactors, respectively. The thermocouples were installed at locations $i \in \{8, 16, 20\}$ in the first reactor and at the following locations of the second reactor: $i \in \{25, 28, 31, 34, 37, 40, 43, 45\}$ along pole 1, $i \in \{24, 27, 30, 33, 36, 39, 42, 45\}$ along pole 2 and $i \in \{23, 26, 29, 32, 35, 38, 41, 45\}$ along pole 3.

We assume that the considered hydrotreatment unit is described well by (1)–(4) and we modeled its pre-processed

Fig. 4. Experiments with real data: The catalyst activity (above) and the temperature differences ΔT (below) reconstructed at different reactor levels by the proposed grid-structured models. Left: Cell with PDE parameterization, middle: MLP-based cell, right: MLP-based cell with regularization. The dotted lines represent the measurements. The ΔT time series are min-max scaled.

TABLE III
PEARSON CORRELATION COEFFICIENT OF TEMPERATURE RECONSTRUCTIONS FOR A REAL REACTOR.

	Train	Valid. set 1	Valid. set 2
Grid-GRU	0.909	0.845	0.860
PDE param	0.797	0.757	0.805
MLP	0.834	0.791	0.798
MLP-reg	0.821	0.785	0.800

data with the proposed grid-structured models. The inputs of the models were inlet concentrations $x_a(z_0, t)$, $x_p(z_0, t)$, inlet temperatures $T(z_0, t)$ and linear velocity $U(t)$. The models were trained to fit temperature measurements $T(z_i, t)$. As the training set, we used the measurements from the first reactor and the sensors attached to the first pole of the second reactor. We used two validation sets which consisted of the temperature measurements from the second and the third poles of the second reactor.

We trained four grid-structured models on the reactor data. The first model was Grid-GRU with 32 neurons trained with learning rates 10^{-3} and 10^{-4} for 15000 iterations each. We did not use the data augmentation that we used for training the same model on the simulated data. This regularization was not needed because the irregular placement of the temperature sensors prevented the type of overfitting that we observed for the simulated data. The other models that we trained were based on the cell defined in (11)–(14). These models were trained with the same hyperparameters as in the experiments with the simulated data. The only difference was that we used data augmentation by applying additive Gaussian noise to all measurements. The noise had zero mean and the standard deviation was 0.002 times the mean of the corrupted variable.

Fig. 4 present the estimated catalyst activity and the fit of the temperature differences ΔT , as defined in (15), which were obtained by the proposed models. One can see that both models capture the drift of the reaction from the inlet towards

Fig. 5. Responses of the trained grid model with the MLP-based cell to changes of the operating conditions. Dashed lines represent temperature differences ΔT predicted by the model for reference operating conditions. The solid lines represent model predictions under changed operating conditions: increase of the inlet temperature by 10% (left), increase of the inlet concentration of P by 25% (middle), increase of the inlet concentration of A by 25% (right).

the outlet of the reactor. The MLP-based model provides a reasonable fit to the data and is comparable with the results obtained with the PDE parametrization. Table III presents the Pearson correlation coefficient of the temperature predictions on the training and two validation sets. These results suggest that all the models can generalize to the validation sets from the same catalyst cycle reasonably well. We did not test the accuracy of the models on a data set with different inputs because of data scarcity.

In Fig. 5, we demonstrate how the models respond to changes of the inputs compared to reference operating conditions. We can see that the increase of the inlet temperature slows down deactivation while the increase of the inlet concentrations of P and A speeds up deactivation, as in (1)–(4).

VII. CONCLUSION

In this paper, we proposed a model that can be used to estimate unobserved states of tubular reactors. The model has a grid-structured architecture which is motivated by the computations performed when solving partial differential equations which describe chemical processes inside such reactors. The model can be viewed as a recurrent network whose computational units are constructed using the prior knowledge about the process dynamics. We demonstrate that the trained model can be used to monitor variables, such as catalyst activity, which can be costly or impossible to measure directly in real-world industrial reactors.

One direction of future research is to use the proposed architecture for monitoring process states using new observations. This can be done, for example, by training the grid-structured model with a generic computational unit such as GRU using data produced by process simulators under randomized scenarios [17].

Due to data scarcity in many industrial domains, there is clear need for more research on combining first-principles models with the data-driven approach. In this paper, we introduced the right inductive bias into a data-driven model by structuring its computational graph, selecting the right computations and by using certain data augmentation schemes. Other promising approaches include learning a combined model containing both a physical component and a data-driven

component [18] or by learning data-driven models that respect conservation laws such as mass or energy balance [19], [20].

ACKNOWLEDGMENT

We would like to thank Akshaya Athwale, Muhammad Emzir, Sakira Hassan, Simo Särkkä, James Kabugo, Teemu Ikonen, Viljami Iso-Markku, Stefan Tötterman, Amir Shirdel, Sanna Laitinen and Kristian Bergman for help in the data gathering process and fruitful discussions. We thank CSC (IT Center for Science, Finland) for computational resources and the Academy of Finland for the support within the Flagship programme: Finnish Center for Artificial Intelligence (FCAI).

REFERENCES

- [1] G. Froment, K. Bischoff, and J. De Wilde, *Chemical Reactor Analysis and Design, 3rd Edition*. John Wiley & Sons, Incorporated, 2010.
- [2] M. E. Davis and R. J. Davis, *Fundamentals of chemical reaction engineering*. Courier Corporation, 2012.
- [3] I. Goodfellow, Y. Bengio, A. Courville, and Y. Bengio, *Deep learning*. MIT Press, 2016, vol. 1.
- [4] T. H. Price and J. B. Butt, “Catalyst poisoning and fixed bed reactor dynamics—II: Adiabatic reactors,” *Chem. Engineering Science*, vol. 32, no. 4, pp. 393–412, 1977.
- [5] K. Cho, B. Van Merriënboer, C. Gulcehre, D. Bahdanau, F. Bougares, H. Schwenk, and Y. Bengio, “Learning phrase representations using RNN encoder-decoder for statistical machine translation,” *arXiv preprint arXiv:1406.1078*, 2014.
- [6] S. Hochreiter and J. Schmidhuber, “Long short-term memory,” *Neural computation*, vol. 9, no. 8, pp. 1735–1780, 1997.
- [7] R. González-García, R. Rico-Martínez, and I. G. Kevrekidis, “Identification of distributed parameter systems: A neural net based approach,” *Computers & chem. engineering*, vol. 22, pp. S965–S968, 1998.
- [8] Y. Lu, A. Zhong, Q. Li, and B. Dong, “Beyond finite layer neural networks: Bridging deep architectures and numerical differential equations,” in *Int. Conf. on Mach. Learn.*, 2018, pp. 3276–3285.
- [9] Z. Long, Y. Lu, X. Ma, and B. Dong, “PDE-net: Learning PDEs from data,” in *Int. Conf. on Mach. Learn.*, 2018, pp. 3208–3216.
- [10] F. d. A. Belbute-Peres, T. Economou, and Z. Kolter, “Combining differentiable PDE solvers and graph neural networks for fluid flow prediction,” in *Int. Conf. on Mach. Learn.*, 2020, pp. 2402–2411.
- [11] R. T. Q. Chen, Y. Rubanova, J. Bettencourt, and D. K. Duvenaud, “Neural ordinary differential equations,” in *Advances in Neural Information Processing Systems*, vol. 31, 2018, p. 6571–6583.
- [12] P. Battaglia, R. Pascanu, M. Lai, D. Rezende, and K. Kavukcuoglu, “Interaction networks for learning about objects, relations and physics,” in *Advances in Neural Information Processing Systems*, vol. 29, 2016, p. 4502–4510.
- [13] H. Deans and L. Lapidus, “A computational model for predicting and correlating the behavior of fixed-bed reactors: I. derivation of model for nonreactive systems,” *AIChE Journal*, vol. 6, no. 4, pp. 656–663, 1960.
- [14] J. Guo, Y. Jiang, and M. H. Al-Dahhan, “Modeling of trickle-bed reactors with exothermic reactions using cell network approach,” *Chem. engineering science*, vol. 63, no. 3, pp. 751–764, 2008.
- [15] J. L. Ba, J. R. Kiros, and G. E. Hinton, “Layer normalization,” *arXiv preprint arXiv:1607.06450*, 2016.
- [16] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” *arXiv preprint arXiv:1412.6980*, 2014.
- [17] J. Tobin, R. Fong, A. Ray, J. Schneider, W. Zaremba, and P. Abbeel, “Domain randomization for transferring deep neural networks from simulation to the real world,” in *IEEE/RSJ int. conf. on intelligent robots and systems*. IEEE, 2017, pp. 23–30.
- [18] V. L. Guen, Y. Yin, J. Dona, I. Ayed, E. de Bézenac, N. Thome, and P. Gallinari, “Augmenting physical models with deep networks for complex dynamics forecasting,” *Int. Conf. on Learn. Represent.*, 2021.
- [19] S. Greydanus, M. Dzamba, and J. Yosinski, “Hamiltonian neural networks,” in *Advances in Neural Information Processing Systems*, vol. 32, 2019.
- [20] M. Cranmer, S. Greydanus, S. Hoyer, P. Battaglia, D. Spergel, and S. Ho, “Lagrangian neural networks,” *arXiv preprint arXiv:2003.04630*, 2020.