INL/CON-23-71412-Revision-0

Adaptive Data-Driven Model Predictive Control for Heat Pipe Microreactors

July 2023

Changing the World's Energy Future

Linyu Lin, Joseph Eugene Oncken, Vivek Agarwal, Benjamin Zastrow

INL/CON-23-71412-Revision-0

Adaptive Data-Driven Model Predictive Control for Heat Pipe Microreactors

Linyu Lin, Joseph Eugene Oncken, Vivek Agarwal, Benjamin Zastrow

July 2023

Idaho National Laboratory
Idaho Falls, Idaho 83415

http://www.inl.gov

Prepared for the
U.S. Department of Energy
Under DOE Idaho Operations Office
Contract DE-AC07-05ID14517

Adaptive Data-Driven Model Predictive Control for Heat Pipe Microreactors

Linyu Lin[1*], Benjamin Zastrow[2], Joseph Oncken[1], Vivek Agarwal[1]

[1]Idaho National Laboratory, Idaho Falls, ID
[2]University of Texas at Austin, Austin, TX

ABSTRACT

To establish a technical basis for self-regulating microreactors, a model predictive control (MPC) system is investigated to proactively respond to anomalies and disturbances in anticipation of potential deviations from operating setpoints. Due to the difficulty of developing a physics-based surrogate model that can accurately match plant data in various operating conditions, machine learning algorithms are used in MPC, which allow for learning from both simulation and operation data, thus efficiently describing the targeted transient with arbitrary accuracy. However, one of the biggest concerns in applying ML algorithms like artificial neural networks (ANNs) is that the predictive capabilities of ANN are limited by training data. If there are gaps between the training and target domain, the accuracy of an ANN can degrade significantly when it is used to predict unseen data. To improve the predictive capability of ANN and enable a confident use of data-driven MPCs outside the training data, this study proposes an adaptive data-driven MPC framework. The system will monitor the discrepancy between plant responses and surrogate predictions, fine-tune the ANN-based surrogate when a large discrepancy is detected, and continue MPC operation with updated surrogates. The framework is demonstrated on a point kinetic model for microreactors. The hyperparameters of the update strategy, including layers to update, error thresholds, learning rate discount, and number of data points used for fine-tuning, are optimized so the simulated microreactor is able to follow changes in setpoint with the smallest of deviations.

Keywords: adaptive model predictive control, neural network

1. INTRODUCTION

Microreactors are a class of very small modular (<50MWe) reactors targeted for special purposes, including remote communities, mining sites, and remote defense bases. Despite various designs, autonomous controls are required for microreactors so the system can govern and regulate itself under different conditions with simplified operations. Such conditions include continual monitoring, prognostics, optimization, and decision-making to improve the operational reliability and resilience in anticipation of potential anomalies, including load variations, plant component degradations, design-basis events, cyber incidents, and other external events [1].

Of the various strategies for autonomous operations, anticipatory control is proposed for enhancing our understanding of how plants or systems of interest must be able to respond to fluctuating loads and adapt to different operating conditions, all with minimal or no human intervention. Unlike reactive or feedback control systems, anticipatory control systems proactively respond to anomalies and disturbances; that is, they afford effective contingency control strategies in anticipation of (not simply in response to) potential impacts to system states [2]. Moreover, considering the compact design and high sensitivity of

* Corresponding author. Email: linyu.lin@inl.gov

microreactors to operating conditions, small disturbances could result in large variations in reactor states. As a result, to accurately manage operating conditions and maintain the microreactor at optimal power levels, an anticipatory control approach is selected for predicting system behaviors and proactively responding before these changes occur.

Based on this principle, model predictive control (MPC) is a technique for implementing anticipatory control strategies. Fundamental to MPC is finding a sequence of control actions that drive the process to the desired setpoints in an optimal way given (1) a discrete-time model that correlates control actions with the controlled variables is available; (2) control actions can be solved from the target values for controlled variables based on the inversed correlations. The principle has been extended to account for mismatches between predictions and targets, multiple input–multiple output systems, and process constraints. A number of surveys on the theory, implementation, and limitations of MPC have been presented in review articles [3], and MPC's most attractive feature for nuclear applications is its ability to handle multiple and complex constraints directly and explicitly. Moreover, MPC is able to deal with complex control problems with multi-inputs, multi-outputs (control actions and state variables), multi-objectives, and various operational constraints. Previous work demonstrated the use of MPC in managing heat pipe temperatures in microreactors [4]. As a result, MPC is well suited for nuclear reactor controls since strict requirements are usually placed on reactor maneuvering and the controller's tracking capabilities, and the control system needs to coordinate and regulate multiple electrical and mechanical components. In addition, increases in computational power continue to make modeling and real-time optimization in MPCs for complex systems ever more feasible.

However, it is apparent from previous descriptions that MPC's capability in tracking time-varying trajectories depend on the accuracy that the process model can approximate the target transients [3]. The first major concern is the use of linear models for nonlinear behaviors, especially for conditions beyond the applicable domain of the linear model. The second concern is the use of a static model to represent the dynamical process in various operating conditions. To confront the problems that arise in controlling nonlinear and time-varying processes, the nonlinear MPC concept was proposed with various techniques. Maiti and Saraf [5] presented an adaptive dynamic matrix control (DMC) where the process is identified in terms of pseudo impulse response coefficients that are updated with time. Dougherty and Cooper developed a multiple model adaptive strategy for a single-loop MPC and also for a multivariable MPC [6]. Their approach mainly relies on combining multiple linear DMCs, whereas the final control action is obtained by interpolating between the individual controller outputs. However, as the authors report, their work focuses on processes that are nonlinear with respect to the operating level, but that are static with respect to time. To address the second concern of MPC, a more practical approach when dealing with time-varying systems is to update the process model in real time. Genceli and Nikolaou [7] proposed a methodology that combines a constrained MPC with simultaneous online model identification for finite impulse response (FIR) models. A constraint was added in the MPC optimization problem that represents the persistent excitation (PE) criterion. Subbaraman and Benosman [8] combined a model-free extremum-seeking (ES) algorithm with MPC to iteratively learn the uncertainties of model parameters with proven convergence and stability. Morton et al. [9] used Deep Variational Koopman model in MPC to continuously infer the posterior distribution of model outcomes.

Although MPC approaches have been improved to better capture nonlinear behaviors or adapt to new operating conditions, nonlinearity and adaptability are rarely considered together. In this article, we propose an adaptive nonlinear MPC configuration with a feedforward neural network (FNN) for predicting transients of selected state variables. The adaptability is implemented by fine-tuning part of the pretrained FNN with new data points when the discrepancies between predicted and measured states are beyond an error limit. The same learning algorithms are used but with discount learning rates. The advantage of fine-tuning is that the weights and biases of the neural network are partially updated so that the FNN can adapt to unseen data points without losing predictive capabilities learned from the original training data. The paper is organized as follows: Section 2 formulates the control problem for the microreactor power transients based on a multiphysics simulation and a simplified lumped parameter model. Section 3 describes the setup of an adaptive MPC with FNN training and fine-tuning. Section 4 presents the numerical results. Section 5 concludes the paper with a summary of the results and a brief discussion on future work.

2. PROBLEM FORMULATIONS

Microreactors are a special class of small reactor systems. Compared to small modular reactors, microreactors are very small in size and target non-conventional nuclear markets. Among various microreactor designs, heat-pipe-cooled (HP-cooled) microreactors use HP elements to remove heat from fission reactions. The system structure is greatly simplified by omitting the main pipeline, circulating pump, and auxiliary equipment. However, compared to large light-water reactors, the power outputs from microreactors contribute to a very large portion of electricity in remote sites. Therefore, the power maneuvering capabilities of microreactors must meet the operation and efficiency requirements of remote electrical grids that have fluctuating electricity demands.

In a generic HP-cooled microreactor design, the fission heat is generated and transferred to the evaporator region of heat pipes through the structural materials with high thermal conductivity. The heat is then transported from the HP evaporator to condenser regions and further removed by an energy conversion system from the HP condenser. Compared to other microreactor designs, HP microreactors enable passive primary coolant, minimal moving parts, and low-pressure operations. However, the compact microreactor design tightly couples different components, and the reactivity is sensitive to dimensional and material changes from the Doppler effect and swelling of the fuel and structures. As a result, the power transients tend to be more complicated than light-water reactors, and a multiphysics model is needed to accurately predict temperature distributions and power transients of microreactors. Since there are not yet any operating HP microreactors, this work starts from Multiphysics Object-Oriented Simulation Environment (MOOSE)-based tools to simulate generic microreactor designs. This work uses similar geometry to the Kilowatt Reactor Using Stirling TechnologY (KRUSTY) [10]. Figure 1(a) shows the scheme of KRUSTY core with safety rods, highly enriched uranium (U-8MO) fuel, eight heat pipes, and reflectors. Compared to the KRUSTY design, as shown in Figure 1(b), this work only models 1/8 of the microreactor core with a single heat pipe [11]. This study uses U-10MO fuel with 93% enrichment and assumes the axial location of the central safety rod is maneuverable and has similar functions to the control rod.

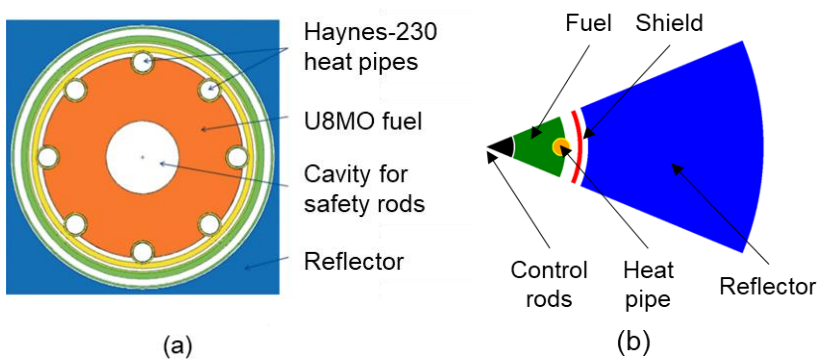

Figure 1: (a) Core scheme of KRUSTY core and (b) simulated 1/8 core in this work.

This work uses Cardinal [12] to couple the neutronic solver OpenMC with a MOOSE heat conduction module to solve the temperature distributions at different power rates, HP temperatures, and control rod insertions. The objective is to understand the temperature feedback and heat capacity of the microreactor core and to inform the design of a point kinetic equation (PKE) with lumped-parameter heat conduction models. Figure 2 shows the OpenMC results for (a) normalized axial power density distributions with different number of particles (b) normalized axial temperature distribution with different control rod positions. The power is normalized with respect to the steady-state power, while the temperature is normalized by the boundary temperatures. Similar to reference [11], Dirichlet boundary conditions are applied at HP locations on the right with a constant HP temperature, while other boundaries are assumed to be adiabatic.

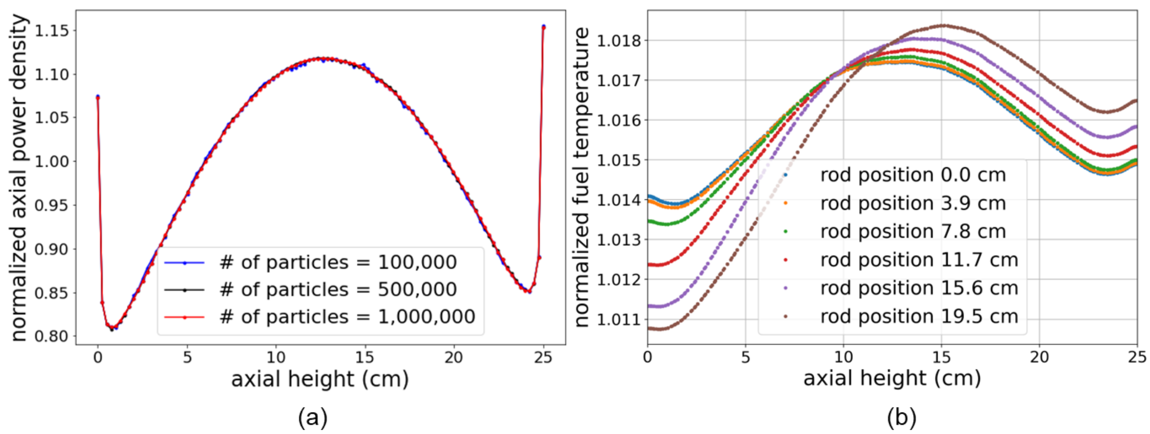

Figure 2: OpenMC results for (a) normalized Axial power density distributions with different number of particles (b) normalized axial temperature distribution by Cardinal. Both plots are made at radial position $\boldsymbol{r = 2.05}$cm, $\boldsymbol{\theta = 0}$.

Although the tool is able to provide a reasonable representation of power and temperature distributions at steady states, transient capabilities are not currently available in OpenMC. Moreover, at the time this work was conducted, higher fidelity simulation tools for power transients of HP-based microreactor were not accessible. As a result, to demonstrate and test an adaptive control strategy, a simplified PKE model with lumped heat conductions was used to simulate the power transients of the microreactor. The reactivity coefficient, heat capacity, and temperatures were derived based on the multiphysics simulation results presented above.

Similar to reference [13], the PKE in this work models the relative density of prompt neutrons n and the precursor concentration of six delayed neutron groups. Meanwhile, heat transfer models with lumped heat-transfer coefficients and average temperatures are built to account for the heat transfer from the fuel to evaporator regions of HPs, from the evaporator to the condenser regions of HPs, and from the condenser to the heat removal system. In addition to external reactivity from control rods, fuel and HP temperature feedbacks are included with temperature feedback coefficients. Table I lists all important model parameters in PKE and lumped heat conduction models. The temperature feedback coefficients for fuel and HP are extracted from reference [10], and all other coefficients are derived from the multiphysics simulations. Since the objective of this work is to develop and demonstrate an adaptive MPC, verification, validation, or development of high-fidelity simulation tools are beyond the scope. This work assumes that the PKE model can represent the target power transient of the microreactor with sufficient accuracy. To demonstrate the nonlinear and adaptive MPC approach, we further assume that a surrogate process model is needed to represent the power transients simulated by PKE.

Table I: Model parameters and initial values of state variables in PKE for microreactor.

Coefficient name	**Value**
Temperature feedback coefficients from fuel	-0.01844
Temperature feedback coefficients from HP	-0.0015
Mass flow rate times heat capacity of secondary heat removal	45.12

Heat transfer coefficient between fuel and evaporator region of HP	17.09
Total heat capacity of the fuel material	26.3
Total heat capacity of the HP	70.5
Reactivity worth of the control rod per unit length	-0.274
Fraction of reactor power deposited in the fuel	93.7%
Effective prompt neutron lifetime	1e-7
Average HP evaporator temperature at steady states	1102
Average fuel temperature at steady states	1277
Initial external reactivity to establish the steady state	2.74
Average HP condenser temperature at steady states	1013
Reactor power rates at steady states	4000
Total delayed neutron fraction	0.0066
Fraction of neutrons from delayed group	[0.00023, 0.0012, 0.0012, 0.0025, 0.0010, 0.00044]
Effective precursor decay constant	[0.013, 0.033, 0.12, 0.30, 0.85, 2.86]

Figure 3 shows reactor power rate transients in response to control rod position changes. A series of power transients is generated based on random combinations of linear ramping, step changes, and sinusoidal movements of control rods. To ensure continuity, the starting position of the next control rod movement is made the same as the ending position of the previous movement. These control rod movements are hypothetical and arbitrary. The goal is not to accurately represent behaviors of the microreactor, but rather to determine if the proposed MPC is able to find optimal control rod movements in extreme cases such that the arbitrary changes in setpoints can be achieved. Such a capability is an important step toward autonomous load-following microreactors.

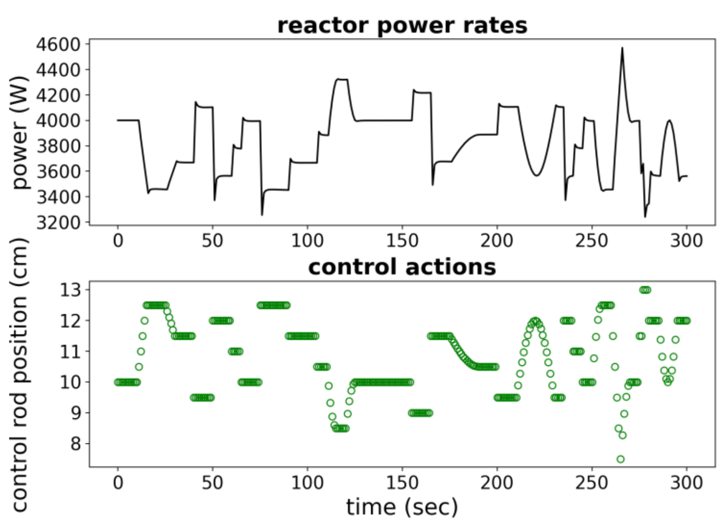

Figure 3: Transient of reactor power rates in response to control rod positions.

3. ADAPTIVE MPC WITH FNN

The essence of MPC is to optimize, over the manipulatable inputs, forecasts of process behaviors subject to equality-inequality constraints. The forecasting is performed using a process model (i.e., a predictor) over a finite time interval of length N. Figure 4 shows the operational scheme of MPC for controlling a microreactor simulator, which is implemented with PKE and mentioned in Section 2. The measurement from the simulator is used as the initial conditions for MPC to identify the optimal control rod positions such that the target trajectories of selected state variables can be achieved. Meanwhile, constraints on control actions and state variables need to be satisfied.

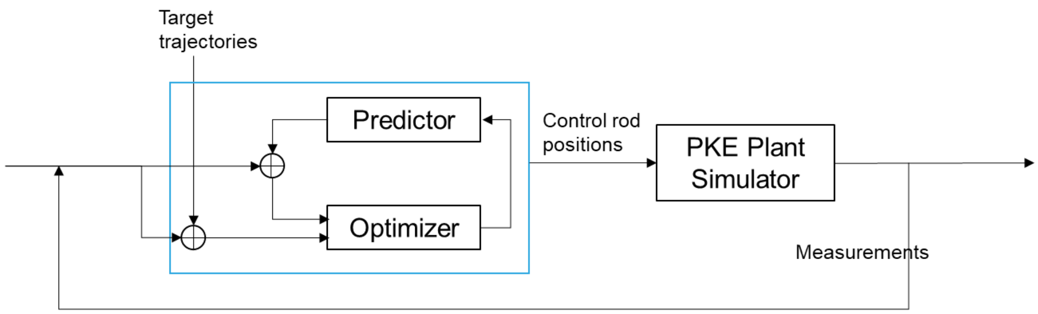

Figure 4: Operational scheme of MPC for a PKE plant simulator for microreactor.

Equation (1) shows a generic form of predictor model f for predicting power $\bar{P}_a$ in response to the axial position z_r of the control rod. This work uses an FNN-based surrogate to approximate the power transient of a microreactor. To evaluate the accuracy of FNN approximations, Equation (2) is used to determine the RMSE ε_P, where P_a and $\bar{P}_a$ represent the actual and approximate reactor powers, respectively.

$$(\bar{P}_a)_{k|j} = f\big((z_r)_{k|j}\big) \tag{1}$$

$$\varepsilon_P = \|\bar{P}_a - P_a\|_2 \tag{2}$$

where j represents the time step, at which point the MPC is activated to find the optimal control action, and $k = 1, \ldots, N$, represents the steps ahead of j that system states are predicted. For convenience, we use subscript $k|j$ to denote k steps ahead of the real-time j, and the initial states of P_a when $k = 0$, $(P_a)_{0|j}$ are obtained from simulation results from the PKE. Equation (3) shows the optimization scheme based on the predictions over a prediction horizon N. The goal is to find a sequence of control-rod positions $(z_r)_{k|j}$ with $k = 1, \ldots, N$ that minimizes the summation J of stage cost function l over the entire prediction horizon N. This work uses a squared error as in (4) with $m = 2$ for deviations between predicted $\bar{P}_a$ and target P_r power rates. Meanwhile, a delta cost penalty is added to the objective function for minor penalties for changing rod positions.

$$J^* = \min_{z_r}\left[\sum_{k=1}^{N} l\big((\bar{P}_a)_{k|j}, (z_r)_{k|j}\big)\right] \tag{3}$$

$$l(\bar{P}_a, z_r) = \|\bar{P}_a - P_r\|_m + \epsilon\|(z_r)_k - (z_r)_{k-1}\|_m \tag{4}$$

subject to

$$(\bar{P}_a)_{k|j} = f\big((z_r)_{k|j}\big)$$
$$Z = \big[(z_r)_{1|j}, \ldots, (z_r)_{N|j}\big] \in \mathbf{Z_i} \text{ for all } i = 1, \ldots, n_{c_z}$$
$$\bar{P} = \big[(\bar{P}_a)_{1|j}, \ldots, (\bar{P}_a)_{N|j}\big] \in \mathbf{P_i} \text{ for all } i = 1, \ldots, n_{c_P}$$
$$(\bar{P}_a)_{0|j} = P_a(j)$$

where J^* stands for the minimized summation of the stage cost function. This work uses a weighted quadratic function to approximate the actual multi-objective, sparse, and nondifferentiable cost function. $\mathbf{Z}_i$ and $\mathbf{P}_i$ represent constraints on the control actions and state variables, respectively. n_{c_z} and n_{c_P} are the numbers of constraints for control actions and state variables, respectively. We use SciPy [14] optimization packages for solving equation (3). Similar to reference [15], this work uses FNN to approximate the correlations between power and control rods. The topology of FNNs consists of a set of neurons linked together in a number of layers. During the design stage, the weights and biases will minimize the empirical risk function with a cost function and a regularization term. In this work, a three-layer FNN with 10 neurons in each layer and the rectified linear unit activation function is trained using the Keras machine-learning package [16]. The training transient is generated by randomly perturbing the control rod positions based on a class of perturbing functions, including linear ramping, step, and sinusoidal functions. The training, validation, and testing data are separated into 80%-10%-10% based on time. For example, with a 100-second transient, the first 80 seconds are used as training data for updating weights and biases of FNN; data from 80–90 seconds are used as validation data, which will stop the training process early if the validation errors stop decreasing for 10 epochs. Data from 90–100 seconds are used to evaluate the accuracy of the trained FNN. If the testing accuracy is not lower than the limit, a new FNN training will be performed with different FNN architecture and training hyperparameters. The initial learning rate is set to be 10^{-3} and reduced by 20% if the training errors are not decreasing for 10 epochs.

The adaptation is achieved by fine-tuning the pre-trained FNN model when the discrepancy ε_P is beyond a limit between the predicted and actual power rates. Next, the measurement data are collected to update a selection of weights and biases in the FNN, while the rest are frozen. Meanwhile, a smaller learning rate is applied to ensure model convergence [17]. The fining-tuning is stopped early if the training errors stop decreasing. The overall procedure for the design and demonstration of adaptive MPC is described as follows:

Algorithm 1: Adaptive MPC with FNN		
Design Procedure: Model development using FNN		
	Generate transient data and split data into training, validation, and testing sets	
	While validation error is decreasing **do**	
		Update FNN weights and biases
	End while	
End Procedure		
Demonstration Procedure: adaptive MPC		
While $t < t_{stop}$ **do**		
	For all time steps k in prediction horizon **do**	
		Solve Equation (1) for $(P_a)_{k\|j}$
	End for	
	Solve Equation (3) for optimal control sequence	
	Apply first action in control sequence and receive measurements	
	If $\varepsilon_P \geq$limit **do**	
		Fine-tune FNN weights and biases
	Else continue	
End while		
End procedure		

The fine tuning starts from an FNN with pretrained weights and biases. Next, only weights and biases in selected hidden layers are updated while the rest are fixed. A smaller learning rate is used to ensure the convergence of FNN training. The ratio of learning rates during fine-tuning and original training is defined as the learning-rate discount. Because of the smaller samples of the dataset, the fine-tuning stops after 50 epochs. Meanwhile, the last 10% of fine-tuning data in time are used in the validation sets for early stopping if the validation errors stop decreasing for 10 steps. To ensure the performance of an adaptive MPC, constants with assigned values in the demonstration stage are optimized with sequential model-based optimization. The constants include discrepancy limits, FNN layers to update, number of collected data for fine-tuning, and learning rate discounts. The hyperparameter optimization is performed with Optuna [18] based on the Tree-structured Parzen Estimator.

4. NUMERICAL RESULTS

This section demonstrates the improved tracking capability of the adaptive MPC with fine-tuning when compared to the MPC without fine-tuning. Figure 5 compares the capabilities of MPC in following the target trajectory with and without FNN fine-tuning. Two fine tunings (labeled as red stars) were performed when deviations in MSE were larger than 300W. The overall RMSE between predictions and measurements was reduced from 382 W to 108 W with an improved tracking capability. With two fine-tunings, the simulation with online updates takes 10 seconds longer than those without any updates.

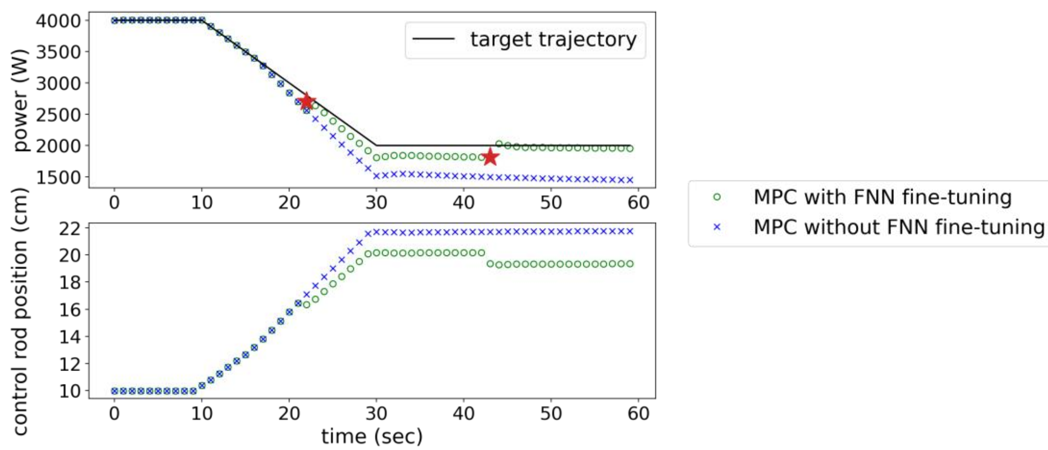

Figure 5: Comparisons of MPC with and without FNN fine-tuning for short transients.

Figure 6 further compares the two MPCs against the baseline MPC with a PKE model instead of surrogates. Table II shows the quantitative results for all three cases, where adaptive MPC shows great improvements in approximating the PKE model and reducing the discrepancy between target and achieved power rates. The final MPC with FNN is able to achieve similar performance to the one with PKE, which is assumed to be the perfect representation of microreactor transients.

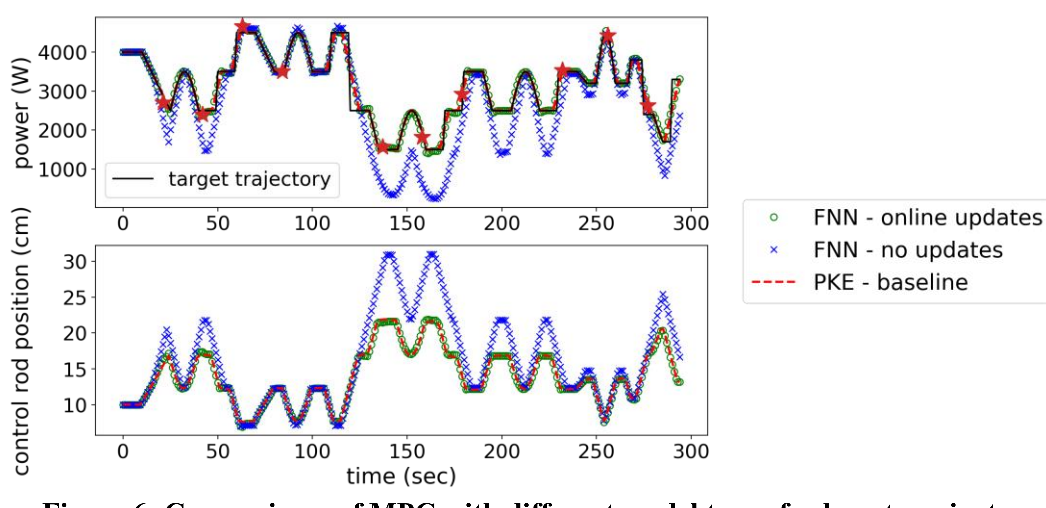

Figure 6: Comparisons of MPC with different model types for long transients.

To improve the performance of online updates, hyperparameter optimization is performed by iteratively drawing samples from uniform distributions of the discrepancy limits, FNN hidden layers, number of collected data for fine-tuning, and learning-rate discount for fine-tuning. The objective is to minimize prediction and tracking errors of adaptive MPC. at the end of the fine-tuning. For a three-layer FNN with

20 neurons, we find that the fine-tuning should be activated when the discrepancy limit is beyond 40 W, while both hidden layers should be updated with optimally 10 data points and a 98% learning rate discount. As shown in Table II, hyperparameter optimizations reduce both the prediction and tracking errors.

Table II: Errors of surrogates in approximating PKE with and without fine tunings. Discrepancy between target and achieved power rates for non-adaptive, adaptive, and baseline MPC with baseline PKE.

	Surrogate models	**RMSE (W)**
Prediction errors	FNN without update	510.0
	FNN with online updates	223.5
	FNN with online updates and hyperparameter optimization	130.5
	PKE baseline	0.0
Discrepancy between target and achieved power rates	FNN without update	649.9
	FNN with online updates	214.7
	FNN with online updates and hyperparameter optimization	178.7
	PKE baseline	168.2

5. CONCLUSIONS

This study demonstrates an adaptive MPC system by fine-tuning an FNN surrogate. The surrogate is used to approximate the microreactor power transient based on which the optimization scheme will search for optimal control actions to meet new targets. Because of gaps between the training data and target scenarios, the discrepancies between surrogate predictions and measurements from the plant simulator could exceed the limit. To reduce the discrepancy, selected weights and biases in the FNN will be updated with new measurement data. Numerical results show that the adaptive MPC is able to significantly reduce the surrogate prediction errors and the discrepancy between target and achieved power rates. The hyperparameters of the update strategy, including layers to update, error thresholds, learning rate discount, and the amount of new data required for fine-tuning, are optimized so the simulated microreactor can follow changes in setpoint with the smallest deviations. One major gap in this manuscript is the validation of simulation models. Future research will account for uncertainties in simulation data and apply the adaptive MPC approach to high-fidelity models for microreactor testbeds.

ACKNOWLEDGMENTS

This work is supported through Idaho National Laboratory (INL) Laboratory Directed Research & Development program under Department of Energy Idaho Operations Office contract no. DE-AC07-05ID14517. This research made use of Idaho National Laboratory computing resources.

REFERENCES

[1] V. Agarwal, K. Manjunatha, A. Gribok, T. Mortenson, H. Bao, R. Reese, T. Ulrish and R. Boring, "Scalable technologies achieving risk-informed condition-based predictive maintenance enhancing the economic performance of operating nuclear power plants (INL/EXT-21-64168)," Idaho National Laboratory, Idaho Falls, ID, 2021.

[2] L. Lin, P. Athe, P. Rouxelin, M. Avramova, A. Gupta, R. Youngblood, J. Lane and N. Dinh, "Digital-twin-based improvements to diagnosis, prognosis, strategy assessment, and discrepancy

checking in a nearly autonomous management and control system," *Annals of Nuclear Energy,* vol. 166, p. 108715, 2022.

[3] M. Schwenzer, M. Ay, T. Bergs and D. Abel, "Review on model predictive control: an engineering perspective," *The International Journal of Advanced Manufacturing Technology,* pp. 1327-1349, 2021.

[4] L. Lin, J. Oncken, V. Agarwal, C. Permann, A. Gribok, T. McJunkin, S. Eggers and R. Boring, "Development and assessment of a model predictive controller enabling anticipatory control strategies for a heat-pipe system," *Progress in Nuclear Energy,* vol. 156, p. 104527, 2023.

[5] S. Maiti and D. Saraf, " Adaptive dynamic matrix control of a distillation column with closed-loop online identification," *Journal of Process Control,* vol. 5, no. 5, pp. 315-327, 1995.

[6] D. Dougherty and D. Cooper, "A practical model adaptive strategy for multivariable model predictive control," *Control Engeering Practice,* vol. 11, p. 649–664, 2003.

[7] Genceli, H., Nikolaou and M., "New approach to constrained predictive control with simultaneous model identification," *AIChE Journal,* vol. 42, no. 10, p. 2857–2868, 1996.

[8] A. Subbaraman and M. Benosman, "Extremum seaking-based iterative learning model predictive control (ESILC-MPC)," *International Federation of Automatic Control,* vol. 49, no. 13, pp. 193-198, 2016.

[9] J. Morton, F. D. Witherden and M. J. Kochenderfer, "Deep Variational Koopman Models: Inferring Koopman Observations for Uncertainty-Aware Dynamics Modeling and Control," *arXiv preprint,* vol. 1902, 2019.

[10] D. I. Poston, M. A. Gibson, T. Godfroy and P. R. McClure, "KRUSTY reactor design," *Nuclear Technology,* vol. 206, pp. 13-30, 2020.

[11] W. Xiao, X. Li, P. Li, T. Zhang and X. Liu, "High-fidelity multi-physics coupling study on advanced heat pipe reactor," *Computer Physics Communications,* vol. 270, p. 108152, 2022.

[12] E. Merzari, H. Yuan, M. Min, D. Shaver, R. Rahaman, P. Shriwise, P. Romano, A. Talamo, Y.-H. Lan, D. Gaston, R. Martineau, P. Fischer and Y. Hassan, "Cardinal: A Lower-Length-Scale Multiphysics Simulator for Pebble-Bed Reactors," *Nuclear Technology,* vol. 207, pp. 1118-1141, 2021.

[13] X. Chen and A. Ray, "On Singular Perturbation of Neutron Point Kinetics in the Dynamic Model of a PWR Nuclear Power Plant," *Sci*, vol. 2, no. 2, p. 36, 2020.

[14] P. Virtanen and e. al., "SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python," *Nature Methods*, vol. 17, no. 3, pp. 261-272, 2020.

[15] L. Lin, P. Athe, P. Rouxelin, M. Avramova, A. Gupta, R. Youngblood, J. Lane and N. Dinh, "Development and assessment of a nearly autonomous management and control system for advanced reactors," *Annals of Nuclear Energy,* vol. 150, p. 107861, 2021.

[16] F. Chollet and e. al., "Keras (https://github.com/fchollet/keras)," GitHub, 2015.

[17] P. Peng and J. Wang, "How to fine-tune deep neural networks in few-shot learning?," *arXiv preprint*, vol. 2012, p. 00204, 2020.

[18] T. Akiba, S. Sano, T. Yanase, T. Ohta and M. Koyama, "Optuna: A Next-generation Hyperparameter Optimization Framework," in *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, Anchorage, AK, 2019.