

A Study on Deep Learning-based 3D Power Profile Prediction Using Ex-core Neutron Detectors

Minhyeok Bang, Dongju Choi, Yunseok Jeong, Junwoo Lee, and Yonghee Kim*

KAIST, 291 Daehak-ro, Yuseong-gu, Daejeon 34141, Republic of Korea

mh_bang@kaist.ac.kr, choidj@kaist.ac.kr, karcice@kaist.ac.kr, junwoolee@kaist.ac.kr,

*Corresponding author: yongheekim@kaist.ac.kr

***Keywords:** 3D Power Distribution, Deep Neural Network, Ex-core Detector, Inverse Problem

1. Introduction

Reconstructing the full 3D power distribution of a nuclear reactor core from sparse, external measurements is a fundamentally challenging inverse problem. Historically, this problem has been solved using in-core detectors. However, this approach carries significant burdens in terms of cost and safety, requiring expensive installation and maintenance as well as reactor pressure vessel penetrations that serve as potential leak paths.

While there have been recent attempts to apply Artificial Intelligence (AI) to this domain, existing studies have predominantly relied on data from in-core detectors [1]. A notable exception is the work by Lin et al. [2], which attempts to forecast power distribution using only ex-core signals. However, its output is limited to a set of 1D axial distributions, failing to reconstruct a full 3D power distribution.

This study aims to accurately predict the full 3D power distribution of the ATOM reactor [3] using only ex-core detector signals. We propose a novel Deep Neural Network (DNN) architecture and investigate the feasibility of using a physics-informed, 'complete' Radial Weighting Factor (RWF) as a proxy for direct in-core measurements. This approach explores whether such a postulated RWF can enhance the model's performance when internal core information is unavailable.

2. Methodologies

2.1. Dataset Generation

A comprehensive and physically realistic dataset was generated by simulating reactor operations at three key burnup stages: Beginning-of-Cycle (BOC), Middle-of-Cycle (MOC), and End-of-Cycle (EOC). For each of these stages, a 63-hour load-following transient was modelled for the ATOM reactor, a 450MWth-class Small Modular Reactor (SMR), utilising our in-house multi-physics PWR simulator code, KANT v1.02.[4] The power manoeuvres for BOC and EOC involved a descent from 100% to 20% power, while the MOC simulation specifically included a more complex 100%-50%-100% power cycle.

A total of 11,342 power profiles were compiled for the training dataset by capturing data at 60-second intervals. This dataset was partitioned, with 90% used for training and 10% random test data set aside for validation (in-domain dataset). Additionally, a distinct validation dataset of 101 data points was generated under the BOC burnup condition by randomly manipulating the control rods to evaluate model robustness (out-of-domain dataset).

The dynamics of the simulated scenarios are illustrated in the accompanying figures. Figure 1 illustrates the reactor power evolution over time. The configuration of the control rod banks is shown in Figure 2, with their corresponding movement during the operation detailed in Figure 3. Finally, Figure 4 provides a statistical overview of the entire 11,342-point dataset, confirming its comprehensiveness and the wide distribution of key core parameters across the three fuel cycles. Figure 5, in turn, illustrates the evolution of its Axial Shape Index (ASI) over time.




Fig. 1. Power evolution of simulation results for BOC, MOC, EOC conditions




Fig. 2. Manoeuvred control rods position of ATOM




Fig. 3. Control rods position evolution of in-domain dataset




Fig. 4. Axial and radial power distributions with ASI of the simulation output results




Fig. 5. ASI evolution of the simulation output data for BOC, MOC, EOC conditions

2.2. Ex-core Detector Response Simulation

For each simulated KANT power profile, the response of the 12 ex-core detectors, arranged as depicted in Fig. 6, was simulated by integrating the power distribution with its corresponding 3D weight map. These weight maps were constructed from the pre-analysed Axial Annealing Functions (AAF) and Radial Weight Functions (RWF) as depicted in Figs. 7 and 8. [4]




Fig. 6. SMART-660 ex-core detector position




Fig. 7. SMART-660 axial annealing function

Truncated Radial Weight Map(Detector B)

1		0	0	0	1.00 E-04	3.80 E-03			
2		0	0	0	0	1.30 E-03	6.28 E-02		
3	0	0	0	0	0	4.00 E-04	1.59 E-02	6.43 E-01	
4	0	0	0	0	0	1.00 E-04	3.50 E-03	2.23 E-01	
5	0	0	0	0	0	0	8.00 E-04	3.88 E-02	
6	0	0	0	0	0	0	2.00 E-04	5.10 E-03	
7	0	0	0	0	0	0	0	9.00 E-04	
8		0	0	0	0	0	0	0	
9			0	0	0	0	0		
	A	B	C	D	E	F	G	H	I

Complete Radial Weight Map(Detector B)

1			1.000 E-07	1.000 E-06	1.000 E-05	1.000 E-04	3.800 E-03		
2			1.000 E-09	1.000 E-08	1.000 E-07	1.000 E-06	1.300 E-05	6.280 E-02	
3	4.000 E-10	4.000 E-09	4.000 E-08	4.000 E-07	4.000 E-06	4.000 E-05	4.000 E-04	1.590 E-02	6.427 E-01
4	1.000 E-10	1.000 E-09	1.000 E-08	1.000 E-07	1.000 E-06	1.000 E-05	1.000 E-04	3.500 E-03	2.233 E-01
5	1.000 E-11	1.000 E-10	1.000 E-09	1.000 E-08	1.000 E-07	1.000 E-06	1.000 E-05	8.000 E-04	3.880 E-02
6	2.000 E-11	2.000 E-10	2.000 E-09	2.000 E-08	2.000 E-07	2.000 E-06	2.000 E-05	2.000 E-04	5.100 E-03
7	2.000 E-12	2.000 E-11	2.000 E-10	2.000 E-09	2.000 E-08	2.000 E-07	2.000 E-06	2.000 E-05	9.000 E-04
8		2.000 E-12	2.000 E-11	2.000 E-10	2.000 E-09	2.000 E-08	2.000 E-07	2.000 E-06	
9			2.000 E-12	2.000 E-11	2.000 E-10	2.000 E-09	2.000 E-08		
	A	B	C	D	E	F	G	H	I

Fig. 8. Axial and radial power distributions with ASI for the simulation input

Ideally, input data from the ATOM reactor should be used for this evaluation, as the RWF and AWF models are tailored for it. However, this study utilises data from the SMART-660 reactor. This decision is justified by the fact that the ATOM and SMART-660 reactors were designed to be identical with same assembly layout and core height. [4] Therefore, it was concluded that using SMART-660 AAF and RWF data would not significantly impede the evaluation of the model's intrinsic performance.[5]

While the truncated model represents a physically correct approach and would typically be used, this study also aimed to investigate a hypothetical scenario: whether the deep learning model could still improve performance by detecting even extremely subtle, complete internal weights. To this end, a complete model was developed for comparison, where internal weights were synthetically derived by applying a 1/10 reduction to adjacent nodes. This hypothesised model serves as an exploratory tool to assess the model's sensitivity to sparse and inferred core information.

2.3. Deep-learning Model Training

A Multi-Layer Perceptron (MLP), a foundational deep learning architecture, was selected as the baseline model for this regression task. The model was trained on the synthetic dataset to learn the complex, non-linear inverse mapping from the 12-channel ex-core signals to the full 3D core power distribution.

The model's architecture consists of an input layer, three hidden layers, and an output layer, with key parameters detailed in Table I. The input layer takes a vector of 12 signals, while the output layer produces a flattened vector of 3,240 nodes corresponding to the 3D power map. The hidden layers progressively increase in size (256, 1024, then 4096 neurons) to gradually expand the model's capacity to reconstruct the high-dimensional output from the sparse input.

Table I: MLP model specification

Parameter	Value / Specification
Total Samples	10,209 training samples, 1,134 in-domain validation samples, 64 out-of-domain validation samples
Input Shape	(12,)
Hidden Layer 1	Dense (256), LeakyReLU ($\alpha=0.2$), Dropout (0.3)
Hidden Layer 2	Dense (1024), LeakyReLU ($\alpha=0.2$), Dropout (0.3)
Hidden Layer 3	Dense (4096), LeakyReLU ($\alpha=0.2$)
Output Shape	(3240,)
Optimiser	Adam
Loss Function	Mean Squared Error (MSE)
Epochs	300

3. Numerical Results

3.1. Model Training and Validation




Fig. 9. Model training and validation loss curves per epoch

The training dynamics, illustrated in the updated Fig. 9, indicate that both models converged effectively within approximately 100 epochs. However, a clear performance difference emerged upon convergence. The truncated model consistently achieved a lower average validation RMSE compared to the complete model, plateauing at a more accurate level. This suggests that for the training dataset, the additional information hypothesized in the complete RWF did not translate into a performance benefit, with the simpler truncated model yielding a more robust solution.

3.2. Validation Result I – In-domain Validation Samples

The model's validation performance was first evaluated on an in-domain validation set, comprising 10% of the whole data that had been randomly excluded from the training set within the same training scenario.




Fig. 10. Radial RMSE distribution of partitioned validation dataset




Fig. 11. Radial RMSE distribution of in-domain validation samples

As illustrated in Fig. 10 and Fig. 11, both the truncated and complete models demonstrated strong predictive performance on the in-domain validation samples. Contrary to expectations, however, the complete model slightly underperformed compared to its truncated counterpart across nearly all metrics.

Fig. 10 shows that while both models exhibited a similar error pattern with a peak at the core centre, the truncated model maintained a consistently lower radial RMSE. This performance advantage is quantified in the summary table of Fig. 11, where the Truncated model achieved a lower overall mean RMSE (1.56% vs. 1.86%) and a lower mean peak power absolute error (0.61% vs. 0.82%) than the complete model.

3.3. Validation Result II – Out-of-domain Validation Samples

Additionally, for a more robust assessment, a completely out-of-domain and novel scenario was considered. This validation dataset comprises 101 data points generated from a simulation of the ATOM reactor under the BOC burnup condition, in which Control Rod 1 (CR1) was forcibly withdrawn to 100%.




Fig. 12. Control rods position evolution of the out-of-domain data




Fig. 13. Power evolution of simulation results of out-of-domain validation samples

As depicted in Fig. 12, the simulation confirms that Control Rod 1 (CR1) is maintained in a fully withdrawn position. This condition induces significant fluctuations in both the core power and the Axial Shape Index (ASI), indicating highly transient reactor behaviour as shown in Fig. 13.




Fig. 14. Radial RMSE distribution of out-of-domain validation samples




Fig. 15. Radial RMSE distribution of out-of-domain validation samples

As shown in Figures 14 and 15, both the truncated and complete models performed similarly on the out-of-domain data, but struggled with significantly higher errors. This was particularly problematic for peak power prediction, where maximum absolute errors reached 23.39% and 24.68% for the truncated and complete models, respectively, revealing a shared limitation in robustness.

4. Conclusions

While our truncated model performed effectively within its domain, the complete RWF unexpectedly offered no additional gains, especially in out-of-domain extrapolation. Moving forward, we plan to significantly boost physical consistency and accuracy by incorporating direct in-core data, like core-top temperatures.

5. Acknowledgements

This work was supported by the Innovative Small Modular Reactor Development Agency grant funded by the Korea Government RS-2024-00405419

REFERENCES

- [1] Li, W., Ding, P., Xia, W., Chen, S., Yu, F., Duan, C., ... & Chen, C. (2022). Artificial neural network reconstructs core power distribution. *Nuclear Engineering and Technology*, 54(2), 617-626.
- [2] Lin, W., Miao, X., Chen, J., Duan, P., Ye, M., Xu, Y., ... & Lu, Y. (2025). Forecasting in-core power distributions in nuclear power plants via a spatial-temporal hierarchical-directed network. *Progress in Nuclear Energy*, 186, 105795.
- [3] Jeong, Y., Choi, D., Oh, T., & Kim, Y. Load-Follow Operation Capability of Soluble Boron-Free Small Modular Reactor ATOM. *Frontiers in Energy Research*, 13, 1639569.
- [4] Oh, Taesuk, et al. "Development and validation of multiphysics PWR core simulator KANT." *Nuclear Engineering and Technology* 55.6 (2023): 2230-2245.
- [5] Roh, G., Kim, K. S., Koo, B. S., Lee, C. C., & Kim, K. Y. (2008). Ex-Core detector response evaluation of the SMART reactor by using the DORT code. *Journal of Nuclear Science and Technology*, 45(sup5), 78-81.