Effect of Surrogate Model on Screening-based Simulated Annealing for OPR-1000 Loading Pattern Optimization

Seongjin Jeong [a], and Hyun Chul Lee [a*]
[a]*School of Mechanical Engineering, Pusan Natl. Univ., 2, Busandaehak-ro 63beon-gil, Geumjeong-gu, Busan, 46241*
[*]*Corresponding author: hyunchul.lee@pusan.ac.kr*

Keywords : loading pattern, OPR-1000, convolutional neural network, vision transformer, simulated annealing

1. Introduction

Loading pattern (LP) optimization is important for improving the fuel-cycle economy of pressurized-water reactors by extending cycle length while satisfying design limits such as the power peaking factor. However, the large combinatorial search space and the high cost of three-dimensional core analyses make direct evaluation of numerous candidate LPs impractical. Simulated annealing (SA) has been widely used for this discrete optimization problem because its probabilistic acceptance rule enables extensive exploration in the early search phase and promotes convergence as the annealing schedule proceeds [1]. To reduce computational costs, surrogate-assisted SA frameworks have been proposed, in which surrogate predictions are used to screen most candidate LPs and 3D core calculations are performed only for uncertain cases, thereby reducing expensive core evaluations while maintaining physics-based verification [2]. In our previous work, this framework was implemented using a pin-wise vision transformer (pin-wise ViT) surrogate and demonstrated that a large portion of candidate LPs can be processed without invoking expensive 3D core calculations, achieving substantial acceleration while preserving feasibility [3].

In this paper, we investigate the effect of surrogate model characteristics on SA behavior and optimization outcomes under an identical optimization setting. Specifically, for the Hanbit Unit 3 Cycle 1 (YGN3 C01) nuclear design report (NDR) core with the fixing fuel assembly (FA) type constraint, surrogate-assisted SA is performed using an assembly-wise convolutional neural network (assembly-wise CNN) surrogate [4] and the pin-wise ViT surrogate [3]. The resulting screening efficiency, computational time, constraint satisfaction, and converged LP characteristics are compared to clarify how the choice of surrogate affects the performance and reliability of the surrogate-assisted SA workflow.

2. Surrogate Models

This section summarizes the two surrogate models that predict cycle length and power peaking factor (max-Fxy) for a given LP and highlights the differences in their input representations and modeling architectures. The detailed definitions of the inputs, network configurations, and implementation settings follow the references [3, 4].

2.1 Assembly-wise convolutional neural network

The assembly-wise CNN uses an assembly-level input representation defined on an 8×8 FA lattice. Each loading pattern is provided as a three-dimensional tensor, where each FA location contains assembly-averaged values including the averaged enrichment, the number of burnable poison (BP) rods, the BP mass fraction, and the assembly-wise burnup. Thus, CNN receives an 8×8 multi-channel map that compactly describes the core at the assembly resolution [4]. The input format is shown in Figure 1.

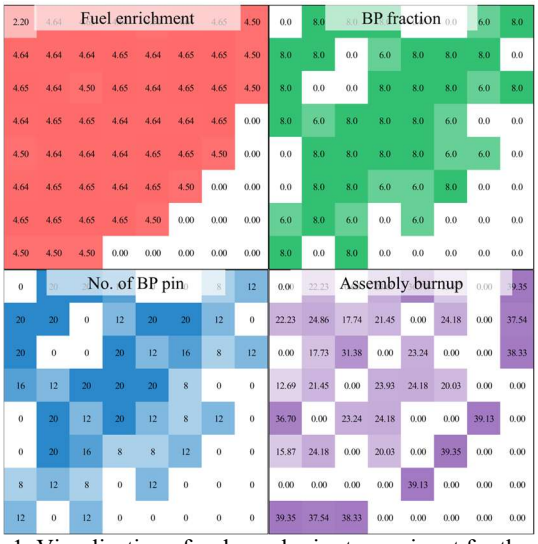

Fig. 1. Visualization of a channel-wise tensor input for the assembly-wise CNN [4].

The assembly-wise CNN adopts a standard pipeline in which multiple convolution blocks learn spatial correlations among neighboring assemblies and progressively build higher-level representations of the LP. The extracted features are then mapped to the target core metrics through a regression head. This structure is well suited to capturing local spatial trends on the assembly map, but its resolution inherently limits the explicit representation of pin-level heterogeneity that can influence power peaking factor behavior. Figure 2 presents a detailed architecture of the assembly-wise CNN for the cycle length evaluation.

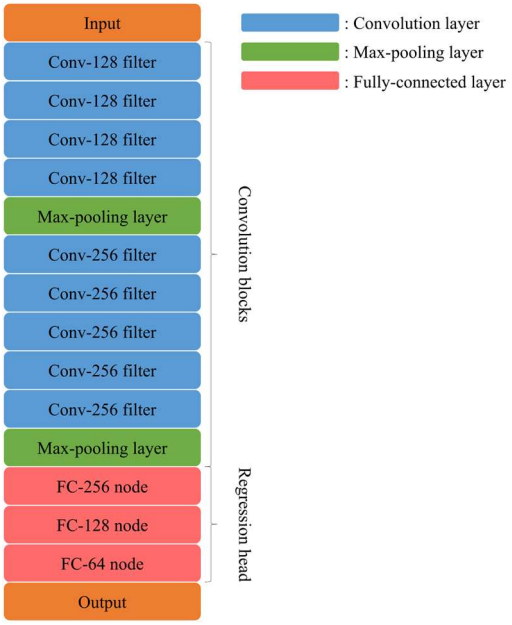

Fig. 2. Architecture of the assembly-wise CNN for cycle-length evaluation [4].

2.2 Pin-wise vision transformer

The pin-wise ViT represents a loading pattern at pin resolution. The input includes pin-level enrichment and pin-level BP mass fraction, which preserve fine spatial heterogeneity within and across assemblies. Burnup is provided as a box-wise-uniform map, where a constant value is assigned within each assembly region to reflect the available burnup information at the box level in the adopted representation. Figure 3 shows the input image format of the pin-wise ViT.

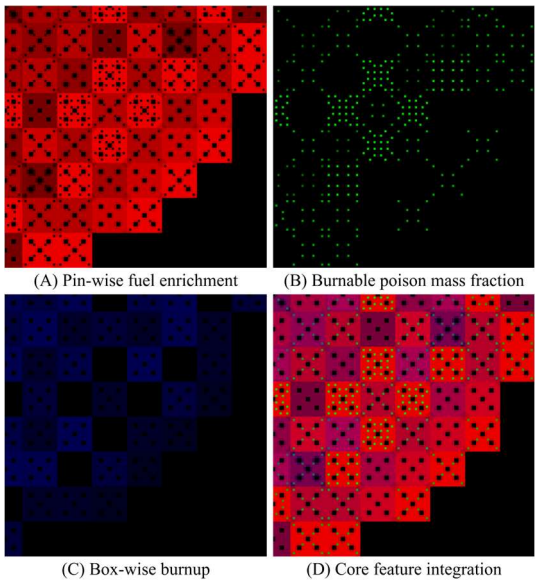

Fig. 3. Visualization of the tensor input for the pin-wise ViT [3].

The pin-wise ViT processes the input by converting the tensor into a sequence of patch embeddings and applying transformer encoder blocks to learn global interactions through attention. This attention-based aggregation allows the model to account for long-range coupling across the core while retaining the spatial detail that can be critical for predicting power peaking factor metrics. The model structure and key design choices follow the previous paper [3]. Figure 4 provides the ViT architecture for cycle length evaluation.

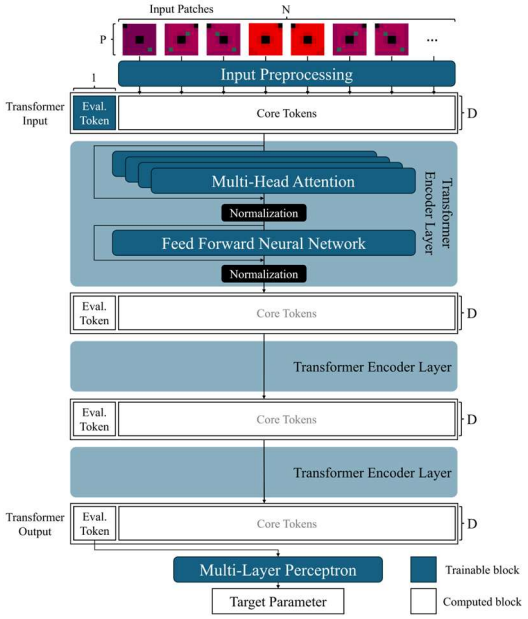

Fig. 4. Architecture of the pin-wise ViT for cycle length evaluation [3].

3. Methods and Results

3.1 Predictive performance of surrogate models

Firstly, both surrogates were trained and evaluated using the same YGN3-based dataset described in Reference [3]. This dataset comprises the random generation (RG), near optimal (NO), and external validation (EV) datasets. The RG dataset consists of randomly shuffled LPs based on the NDR core, while the NO dataset contains near-optimal LPs collected by SA runs using a pre-trained pin-wise ViT with only the RG dataset. To assess generalization to unseen FA types, the EV dataset was generated by replacing several existing FA types with new types having different enrichment levels, numbers of BP rods, or BP fractions, and then applying the resulting modified FA-type list to selected LPs from the original RG and NO datasets.

However, because the EV dataset was largely derived from the RG and NO datasets, it still had a potential limitation in terms of LP diversity. LPs outside both the RG dataset and the NO region can appear as candidate LPs during the pre-SA process, which means that the EV dataset may not accurately represent the predictive error of the surrogate for LPs encountered during the pre-SA process. Figure 5 compares the surrogate predictions with the corresponding 3D core calculation results for the

power peaking factor of LPs generated during the pre-SA process. As shown in the figure, the pin-wise ViT produced similar trends for the EV dataset and LPs generated during the pre-SA process, whereas the assembly-wise CNN showed large prediction bias for the pre-SA LPs that were not observed in the EV dataset. This result likely arises because the assembly-wise CNN uses relatively low-resolution assembly-level inputs, making it more difficult for the model to capture the contextual information needed for power peaking factor prediction.

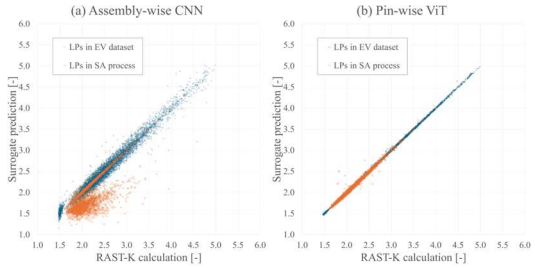

Fig. 5. Comparison of surrogate predictions and 3D core calculation results for the EV dataset and LPs generated during the pre-SA process.

Because such large prediction bias can interfere with the SA process, additional training was conducted for the assembly-wise CNN using 90% of the candidate LPs generated during the pre-SA process for YGN3 C01, and the remaining 10% were used for validation. Figure 6 compares the predictions with the 3D core calculation results for the remaining 10% of LPs before and after the additional training. After this training, the surrogate no longer showed anomalous deviations.

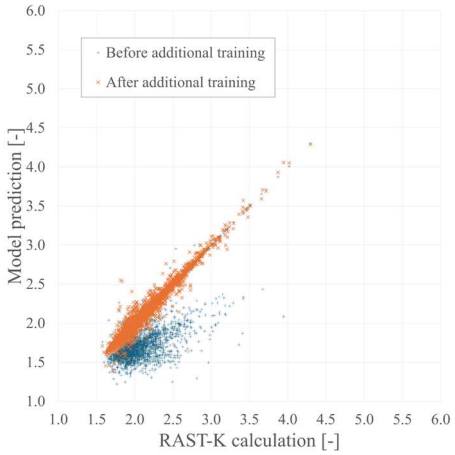

Fig. 6. Comparison of assembly-wise CNN predictions with 3D core calculation results for LPs generated during the pre-SA process.

The final predictive performance of the surrogate models was evaluated on the EV dataset and LPs generated during the pre-SA process, using the Root Mean Square Relative Error (RMSRE) and the Maximum Absolute Relative Error (MaxARE). As summarized in Table I, even without additional training, the pin-wise ViT achieved consistently better predictive performance than the assembly-wise CNN, with RMSRE values that were approximately five times lower.

Table I: Error metrics for assembly-wise CNN and pin-wise ViT on the EV dataset and LPs generated during the pre-SA process.

Error metric	Cycle length (%)		Power peaking factor (%)	
	CNN	ViT	CNN	ViT
RMSRE	0.94	0.16	8.19	1.57
MaxARE	10.0	0.90	40.5	37.6

3.2 Surrogate-assisted simulated annealing

The fixing-type option is adopted to isolate the effect of surrogate selection on the SA search. Under the fixing-type constraint, the total number of FAs in each FA type is held constant throughout the optimization, and candidate LPs are generated only by positional shuffling. Accordingly, the search space is restricted to permutations of a fixed FA inventory, which eliminates composition changes and allows a direct comparison of surrogate-driven differences in SA behavior and outcomes.

A multi-objective adaptive SA algorithm [3] is used to optimize the LP while satisfying the cycle length and power peaking factor constraints. Candidate LPs are generated using the predefined shuffling move set and are first evaluated by the surrogate. Clearly unfavorable candidates are screened out, whereas three-dimensional core calculations are performed only for uncertain candidates that may influence the acceptance decision. The same optimization settings are applied to both surrogates.

The target core is the YGN3 C01 NDR core [5]. Table II provides the detailed FA specifications used in the optimization, including the relevant enrichment and BP configuration information as described in the nuclear design report (NDR) of the core. For each surrogate, ten independent SA runs are conducted under identical settings and executed in parallel to evaluate run-to-run variability and to compare reliability in reaching improved patterns relative to the reference design.

Table II: Fuel assembly specifications for the Hanbit Unit 3 Cycle 1 loading pattern [5].

FA Type	Fuel Enrichment [wt.% U-235]		No. Rods		Burnable Poison Fraction [wt.% Gd2O3]
	Normal	Zoned	Zoned	BP	
A0	1.30	-	-	-	-
B0	2.37	-	-	-	-
B1	2.36	1.30	52	8	4.0
B2	2.37	-	52	4	4.0
C0	2.87	2.35	-	-	-
C1	2.87	2.36	52	8	4.0
D0	3.35	2.87	-	-	-
D1	3.36	2.85	52	8	4.0
D2	3.35	2.87	100	8	4.0

3.3 Optimal loading patterns by the assembly-wise CNN

The assembly-wise CNN is first applied to surrogate-assisted SA under the fixing-type constraint. Table III summarizes the key statistics of the SA process and the final optimization outcomes. Across the ten runs, approximately 89–91% of candidate patterns are decided using surrogate screening, indicating that the surrogate reduces the number of 3D core calculations substantially. Nevertheless, the remaining fraction of uncertain candidates still requires a considerable number of 3D core calculations, and the total runtime remains on the order of 150–300 h per run under the applied computing configuration.

Table III: Results of simulated annealing using the assembly-wise CNN.

SA run	3D core calculations	Assembly-wise CNN evaluations	Screening efficiency	CPU time (h)	Estimated CPU time without screening (h)	Cycle length (EFPDs)	Power peaking factor (-)
1	2,902	28,012	90.6	242.6	2576.2	374.5	1.5254
2	2,754	22,358	89.0	230.1	2092.7	374.8	**1.5423**
3	3,135	32,187	91.1	262.1	2943.5	373.5	1.5396
4	2,355	20,292	89.6	196.8	1887.3	373.8	1.5378
5	1,795	17,845	90.9	150.1	1636.7	373.7	1.5390
6	3,321	29,061	89.7	277.6	2698.5	374.5	1.5388
7	2,825	27,604	90.7	236.2	2535.8	374.8	1.5402
8	3,515	32,023	90.1	293.8	2961.5	374.6	**1.5468**
9	2,252	18,344	89.1	188.2	1716.3	374.5	1.5388
10	2,675	23,378	89.7	223.6	2171.1	373.5	1.5391
Avg.	2,752.9	25,110.4	90.1	230.1	2,321.9	374.4	1.5388
NDR core						373.2	1.5412

In terms of optimization outcomes, two of the ten assembly-wise CNN-assisted runs failed to identify LPs that improved upon the reference design relative to the NDR baseline. This behavior indicates that, despite meaningful computational savings from screening, the assembly-wise CNN does not consistently support stable convergence to improved solutions under the same SA setting. Figure 7 compares the best loading pattern obtained from the CNN-assisted runs with the NDR LP to illustrate representative differences in spatial distribution. As a supplementary check on the stability of the optimized LP, the most positive moderator temperature coefficient (MTC) over the cycle was also evaluated. Because the YGN3 C01 core contains a relatively large amount of BP, the critical boron concentration reaches its maximum during the cycle rather than at the beginning of cycle, and the most positive MTC appears at the same burnup. For both the NDR core and the optimized LP core, the most positive MTC occurred at 54.3 EFPDs and was calculated as -10.3 and -9.7 pcm/K, respectively. Although MTC was not included in the LP optimization objective, it remains sufficiently below zero, indicating that the optimized LP maintains acceptable MTC behavior.

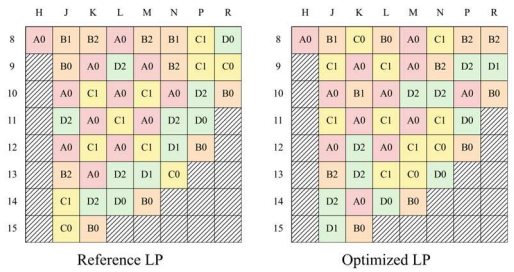

Fig. 7. Reference LP and the representative optimal LP using the assembly-wise CNN.

3.4 Optimal loading patterns by the pin-wise ViT

The pin-wise ViT is also applied to the same surrogate-assisted SA workflow, using the identical fixing-type constraint, move generation, and SA schedule. Table IV reports the corresponding SA statistics and optimization outcomes. In contrast to the assembly-wise CNN case, the pin-wise ViT-assisted runs achieve screening efficiency of approximately 99.9%, meaning that only a small fraction of candidates require 3D core calculation. As a result, the CPU time is reduced to approximately 2–4 h per run under the same computational environment, indicating a large reduction in the number of 3D core calculations relative to the assembly-wise CNN-assisted SA.

Table IV: Results of simulated annealing using the pin-wise ViT.

SA run	3D core calculations	Pin-wise ViT evaluations	Screening efficiency	CPU time (h)	Estimated CPU time without screening (h)	Cycle length (EFPDs)	Power peaking factor (-)
1	24	22,591	99.9	2.6	1,884.6	374.4	1.5347
2	31	23,700	99.9	3.2	1,977.6	374.7	1.5276
3	18	21,936	99.9	2.1	1,829.5	374.7	1.5276
4	31	23,485	99.9	3.2	1,959.7	375.1	1.4965
5	31	23,473	99.9	3.2	1,958.7	375.1	1.4965
6	29	23,389	99.9	3.1	1,951.5	373.9	1.5351
7	33	23,530	99.9	3.4	1,963.6	374.4	1.5347
8	33	23,583	99.9	3.4	1,968.0	374.4	1.5347
9	30	23,425	99.9	3.2	1,954.6	373.7	1.5390
10	20	22,149	99.9	2.3	1,847.4	374.7	1.5276
Avg.	28.0	23,126.1	99.9	3.0	1,929.5	374.5	1.5254
NDR core						373.2	1.5412

The pin-wise ViT-assisted optimization also shows improved reliability in terms of the final outcomes. All ten runs successfully produce LPs that improve upon the NDR baseline while satisfying the power peaking factor constraint under the same optimization setting. Figure 8 compares the best LP among the pin-wise ViT-assisted results with the NDR LP. The MTC of this optimized LP

was also evaluated. The most positive MTC was -10.9 pcm/K at 54.3 EFPDs, the same burnup point as in the NDR core, indicating that this LP also maintains sufficiently negative MTC behavior.

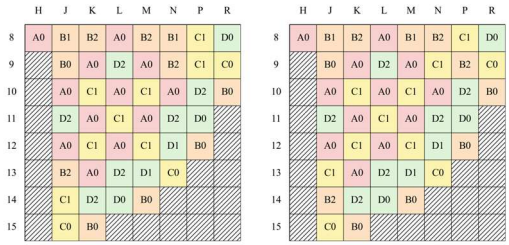

Fig. 8. Reference LP and the representative optimal LP using the pin-wise ViT.

3.5 Surrogate-driven differences in the optimization

Although additional training improved the accuracy of the assembly-wise CNN for diverse LPs, its precision in predicting the power peaking factor remained lower than that of the pin-wise ViT. This lower precision widened the ±3σ uncertainty band and caused more LPs to be passed to full three-dimensional core calculations. This tendency became more pronounced in the later stages, where the search increasingly converged toward stable LPs. As a result, SA using the pin-wise ViT completed the search and converged within only a few hours, whereas SA using the assembly-wise CNN required approximately 150 to 300 h.

In addition, the remaining overestimation and underestimation for a small number of LPs still functioned as artificial barriers or wells during the SA process, thereby restricting moves toward potentially improved LPs or, in some cases, directing the search toward LPs that were less favorable for improvement. Consequently, as shown in the preceding tables, the assembly-wise CNN required more stages and a larger total number of LP evaluations than the pin-wise ViT. Moreover, in two of the ten SA runs, the final LP failed to satisfy the power peaking factor target. These results indicate that the lower input resolution of the assembly-wise CNN, and the resulting limitations in power-peaking-factor prediction accuracy and precision, had a substantial adverse effect on both the efficiency and the overall optimization performance of SA compared with the pin-wise ViT.

4. Conclusions

This study compared assembly-wise CNN and pin-wise ViT for surrogate-assisted SA in the fixing-type LP optimization of the YGN3 C01 core. Both surrogates were initially trained using the same RG, NO, and EV datasets. However, because the EV dataset was largely derived from RG and NO, it had a limitation as a validation dataset for LPs likely to appear during the SA process. Therefore, additional LPs were constructed through a pre-SA process to examine surrogate behavior beyond the EV dataset. In this analysis, the pin-wise ViT showed stable predictions, whereas the assembly-wise CNN showed large deviations. The assembly-wise CNN was therefore additionally trained using LPs obtained from the pre-SA process, while a portion was reserved for validation. Even without additional training on the pre-SA LPs, the pin-wise ViT showed consistently better predictive performance than the assembly-wise CNN.

In ten SA runs, the retrained assembly-wise CNN achieved a screening efficiency of 89–91% and required 150–300 h. All runs satisfied the cycle length constraint, but two failed to satisfy the power peaking factor constraint. In contrast, the pin-wise ViT achieved a screening efficiency of about 99.9%, required about 2–4 h, and produced LPs that exceeded the NDR baseline in all ten runs. These results indicate that, for the present problem, the pin-wise ViT was more effective than the retrained assembly-wise CNN in terms of screening efficiency, computational time, and optimization reliability.

Future work will focus on two main directions. One is to analyze the cause of the limited optimization reliability observed in the assembly-wise CNN. The other is to accelerate the optimization process by replacing simulated annealing with reinforcement learning.

Acknowledgements

This research was supported by the National Research Foundation of Korea (NRF) grants funded by the Korea Government (MSIT) (No. RS-2024-00436693) and the Korea Institute of Energy Technology Evaluation and Planning (KETEP) and the Ministry of Trade, Industry & Energy (MOTIE) of the Republic of Korea (No. 20214000000410).

REFERENCES

[1] D. J. Kropaczek and P. J. Turinsky, In-Core Nuclear Fuel Management Optimization for Pressurized Water Reactor Using Simulated Annealing, Nuclear Technology, 95(1), 9-32 (1991).

[2] T.K. Park, H.G. Joo, C.H. Kim, and H.C. Lee, Multiobjective Loading Pattern Optimization by Simulated Annealing Employing Discontinuous Penalty Function and Screening Technique, Nuclear Science and Engineering, 162(2), 134–147 (2009).

[3] S. Jeong and H. C. Lee, Loading Pattern Optimization for OPR-1000 by Adaptive Simulated Annealing with a Screening Technique using Pin-wise Vision Transformer, Reactor physics Asia conference 2025, (2025).

[4] H. Jang, H.C. Lee, and H.C. Shin, Refinement of convolutional neural network for neutronic design parameter prediction of a loading pattern, Reactor physics Asia conference 2019, 175–178 (2019).

[5] S.K. Lee, et al., Nuclear Design Report for Yonggwang Nuclear Power Plant Unit 3 Cycle 1, Korea Atomic Energy Research Institute (1995).