Loading Pattern Optimization for OPR-1000 by Simulated Annealing with a Screening Technique using Pin-wise Vision Transformer Model

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, screening technique, optimization, OPR-1000, vision transformer

1. Introduction

For decades, research has focused on optimizing the core fuel loading pattern (LP) to extend reactor cycle length while enhancing power stability and fuel efficiency. Early methods such as genetic algorithms evaluated numerous random LPs to find improved configurations [1], but this brute-force approach was inefficient. To overcome this, Simulated Annealing (SA) was introduced [2], significantly reducing the number of evaluations needed to find optimal LPs.

To further enhance efficiency, a screening technique was applied to LP optimization [3], using 2D core deterministic calculations to filter out unfavorable or sufficiently optimal LPs early in the SA process. However, even 2D calculations introduced notable computational overhead. To address this, Convolutional Neural Networks (CNNs) were adopted [4], offering faster evaluations but only at the assembly-wise level, losing crucial pin-wise detail.

To resolve this, a Vision Transformer (ViT)-based LP evaluation model [5] was developed to predict key reactor performance metrics using quarter-core fuel rod distributions, preserving pin-level accuracy while maintaining efficiency. Building on this, the present study proposes a multi-cycle training methodology for the ViT-based LP evaluation model, improving its generalization and predictive accuracy across diverse LP configurations. The fine-tuned model was then integrated into the SA screening process and applied to optimize LPs for the first and second cycles of Hanbit (Yonggwang) Unit 3 [6], a Korean OPR-1000 reactor.

2. Pin-wise Vision Transformer model

2.1. Model architecture

The LP evaluation model consists of two types: the cycle length evaluation model and the peaking factor evaluation model. Each model takes a three-dimensional tensor as a 120x120x3 tensor as input, representing a quarter-core divided into 120x120 fuel rods. For each rod, the input features include the enrichment, the burnable poison (BP) mass fraction, and the burnup value of the quarter-assembly to which the rod belongs. The values for areas outside the core and water holes are set to zero.

Figure 1 presents an example of the input data visualized as an image. In the image, red represents enrichment, green indicates the BP mass fraction, and blue corresponds to the burnup value of the quarter-assembly. Each rod feature is normalized so that the maximum value is scaled to 256, allowing for effective visualization of the input data.

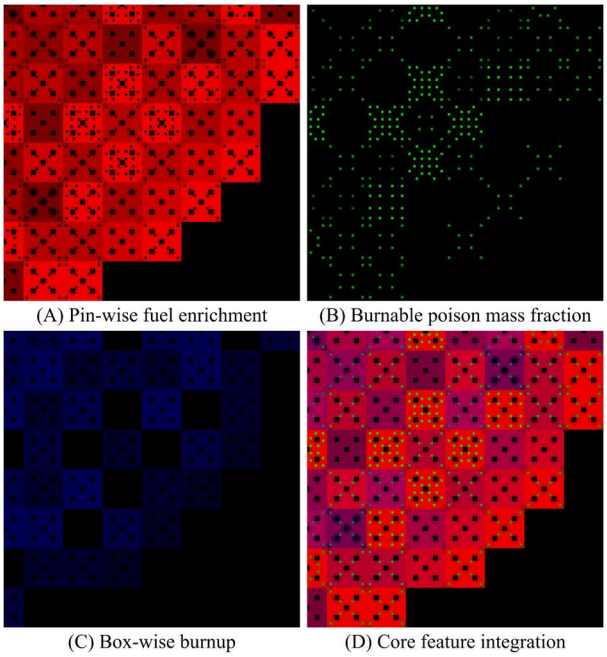

Fig. 1. Visualization of a tensor input for loading pattern evaluation models [5].

The proposed model structure follows a similar design to that presented in a previous study [5], utilizing the ViT framework to effectively analyze two-dimensional quarter-core data. Figure 2 illustrates the overall structure of the LP evaluation model, depicting its key components and processing flow, while Table 1 outlines the essential hyperparameters that define each model configuration.

To adapt the Transformer architecture to the input data, the 120×120×3 tensor is divided into size P×P×3 patches (where P denotes the patch size), each of which is transformed into a 1D tensor and projected onto a hidden dimension D using a dense matrix. The number of patches, N, is given by $(120/P)^2$. The resulting $N \times D$ input tensor is augmented with trainable $N \times D$ positional embeddings, and an additional D-dimensional learnable token is appended for predicting cycle length or peaking factor, resulting in a final $(N + 1) \times D$ tensor.

The Transformer model consists of L stacked layers, each containing a Multi-Head Attention (MHA) and a Feed Forward Neural Network (FFNN). MHA assigns importance weights to tokens using H parallel attention heads, which extract diverse dependencies from the input sequence. The outputs from these heads are concatenated, transformed, and multiplied with token representations to emphasize critical information. FFNN refines extracted features through two fully connected layers, expanding the feature dimension to D_{FFNN} before projecting it back to D. The Gaussian Error Linear Unit (GELU) activation function is used to enhance nonlinearity and adaptability. Residual connections are applied in both MHA and FFNN to stabilize training and facilitate deep learning.

After processing through the Transformer layers, the appended prediction token encodes the necessary information for evaluation. This token is passed through a Multi-Layer Perceptron (MLP) consisting of an input layer, two hidden layers (dimension D_{MLP}), and an output layer, with GELU activation applied in the hidden layers. The final output layer generates the predicted cycle length or peaking factor, providing an accurate estimation of core performance parameters.

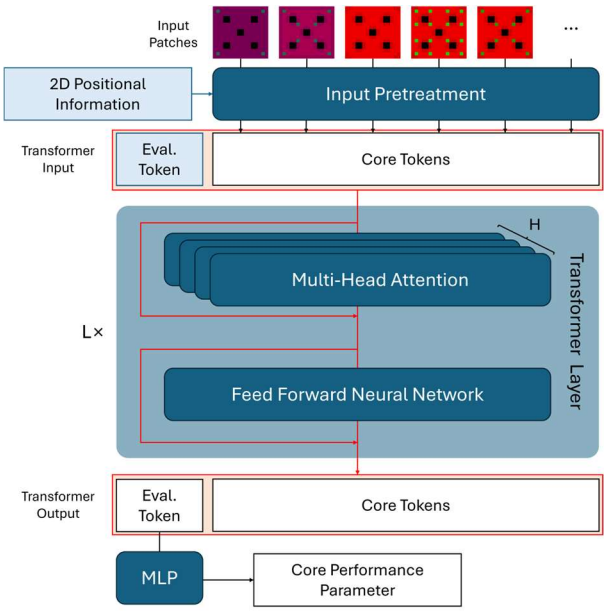

Fig. 2. Pin-wise Vision Transformer model structure.

Table I: Hyperparameters of the pin-wise Vision Transformer models.

Type	Cycle length	Peaking factor
Hidden dimension	256	256
Patch size	8	2
No. of layers	3	5
No. of attention heads	4	4
Dimension of FFNN layer	512	512
Dimension of MLP layer	64/32	64/32

2.2. Model training method

2.2.1. Random LP dataset

To enable the model to analyze multi-cycle LPs, it is essential to train the model with a diverse set of LPs within the required scope. To achieve this, we generated LPs based on Hanbit Unit 3 Cycle 1 and Cycle 2 from NDRs, while Cycles 3 to 16 were derived from an equilibrium LP design paper [6, 7]. Table 2 presents the fuel assembly types used in random LP generation. In the table, assembly types from A0 to D2 represent fresh fuel assemblies for Cycle 1, E0 to E2 for Cycle 2, FC to F6 for Cycles 3 to 10, and GC to GH6 for Cycles 11 to 16. Figure 3 represents types of enrichment zoning pattern and burnable poison arrangement for each assembly.

Table II: Fuel assemblies for loading pattern from Hanbit unit 3 [6, 7].

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
E0	4.08	3.61	-	-	-
E1	4.08	3.61	52	8	6.0
E2	3.60	3.11	52	8	6.0
FC	2.20	-	-	-	-
F0	4.65	4.10	-	-	-
F1	4.65	4.10	52	4	6.0
F2	4.65	4.10	52	8	6.0
F3	4.65	4.10	52	12	6.0
F4	4.65	4.10	52	16	8.0
F5	4.65	4.10	52	20	8.0
F6	4.65	4.10	52	24	8.0
GC	2.90	-	-	-	-
G5	4.75	4.45	52	20	8.0
GH1	4.95	4.45	52	4	8.0
GH2	4.95	4.45	52	8	8.0
GH3	4.95	4.45	52	12	8.0
GH4	4.95	4.45	52	16	8.0
GH5	4.95	4.45	52	20	8.0
GH6	4.95	4.45	52	24	8.0

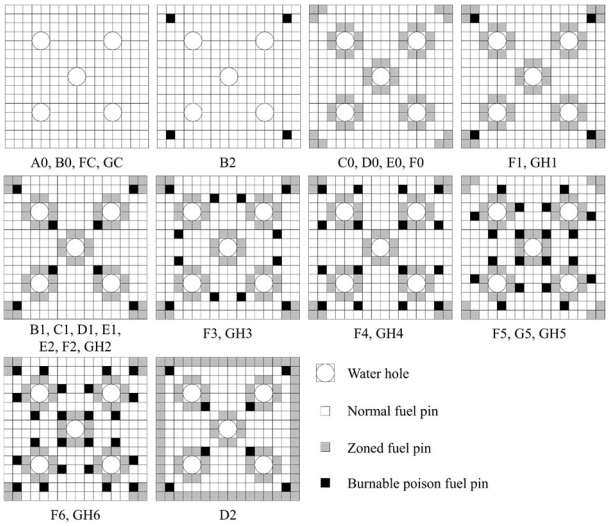

Fig. 3. Enrichment zoning pattern and burnable poison arrangement [6].

For Cycle 1, 5,000 LPs were generated by randomly swapping or replacing fuel assemblies within the NDR LP or designated regions (A0 to D2) used in Cycle 1. Each LP maintained octant symmetry, and as all assemblies were fresh fuel, burnup was set to 0.

For Cycle 2, LPs were generated similarly, but burnup values were assigned to fuel assemblies by dividing them into 2×2 regions, with burnup randomly set within 5–30 MWD/kgU range, with an additional variation of ±5 MWD/kgU. Fresh fuel assemblies retained 0 burnup. 25,000 LPs were generated for Cycle 2.

For Cycles 3 to 10, LPs were derived from 18-month equilibrium cores, using FC to F6 as fresh fuel. Fuel from previous cycles retained 5–30 MWD/kgU (±5 MWD/kgU), while two-cycle-old fuel had 10–60 MWD/kgU (±10 MWD/kgU). 25,000 LPs were generated per cycle for Cycles 3 and 4, while 10,000 LPs were generated per cycle for Cycles 5 to 10 due to similar assembly types.

For Cycles 11 to 16, single-zone LPs with 93 fresh fuel assemblies from 24-month equilibrium cores were used as references. The same approach as Cycles 3–10 was followed, but cycle lengths were longer, so burnup values were adjusted: fuel from the previous cycle had 5–40 MWD/kgU (±5 MWD/kgU), and two-cycle-old fuel had 10–70 MWD/kgU (±10 MWD/kgU). Other aspects of LP generation remained consistent with Cycles 3–10.

A total of 230,000 generated LPs were evaluated using STREAM/RAST-K (ST/RK), a two-step deterministic code developed by UNIST [8], to determine the cycle length and the peaking factor for each LP. If the peaking factor of any LP exceeded 5.0, the LP was regenerated to ensure that the peaking factor remained at or below this threshold.

2.2.2. Near-optimal LP dataset

Random LP generation methods rarely produce near-optimal LPs, making it difficult for the LP evaluation model to accurately predict LPs in the optimal region. To address this issue, we utilize the pre-trained model to perform Simulated Annealing (SA) and collect optimal LP data for improved model accuracy. SA is performed 2,000 times for each cycle from Cycle 1 to Cycle 16. The initial LPs are selected from NDR LPs and those presented in the equilibrium core study. The burnup values of burnt fuel assemblies are obtained from previous cycle calculations.

2.2.3. Method for model training

The dataset is divided into three subsets. The first is the training dataset, which is used to optimize the model's parameters by making predictions based on LPs and minimizing errors through comparison with labeled data. The second is the validation dataset, which is separate from the training dataset and is used to assess the model's error after each epoch. This dataset is not involved in training but serves to evaluate how well the model generalizes to unseen data. Finally, the test dataset is separate from both the training and validation datasets and is used to assess the model's final prediction accuracy after training is complete. The training, validation, and test datasets were split into an 8:1:1 ratio. Figures 4 and 5 illustrate the distribution of cycle lengths and peaking factors for each subset of the randomly generated LPs.

Figure 6 presents the distributions of cycle length in the training, validation, and test datasets derived from the near-optimal LP dataset. The distribution patterns are consistent across all three datasets, ensuring that the model is trained and evaluated on a representative dataset. Multiple peaks indicating the presence of distinct LP categories. This diverse distribution enables the model to generalize effectively across different LP configurations.

Figure 7 shows the distributions of peaking factor for training, validation, and test datasets. The distributions exhibit a sharp peak around 1.5 to 1.6, reflecting the dominance of near-optimal LPs with peaking factors in this range. The validation and test datasets follow a similar distribution to the training dataset, ensuring consistency in model evaluation. Compared to the cycle length distribution, the peaking factor distribution appears more concentrated, suggesting that most LPs exhibit peaking factors within a narrow range.

The model training process follows these steps: In each training step, a batch of 32 samples from the training dataset is processed to compute the model’s output. This output is then compared with the ST/RK results to calculate the Root Mean Square (RMS) error. Based on this error, the backpropagation algorithm is applied to optimize the model’s parameters. This process is repeated until the model has been trained on the entire training dataset, marking the completion of one epoch.

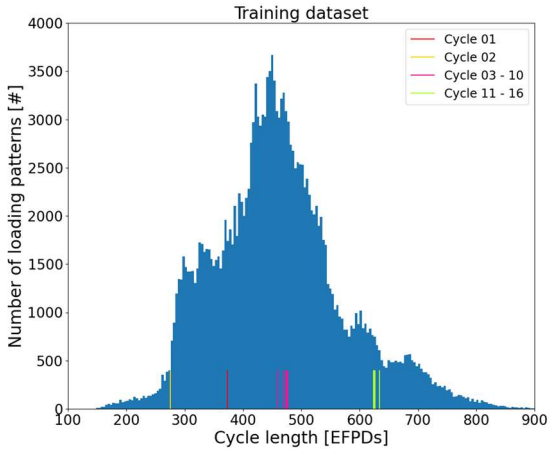

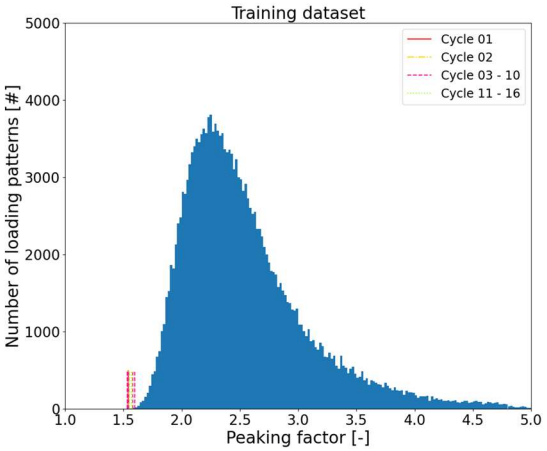

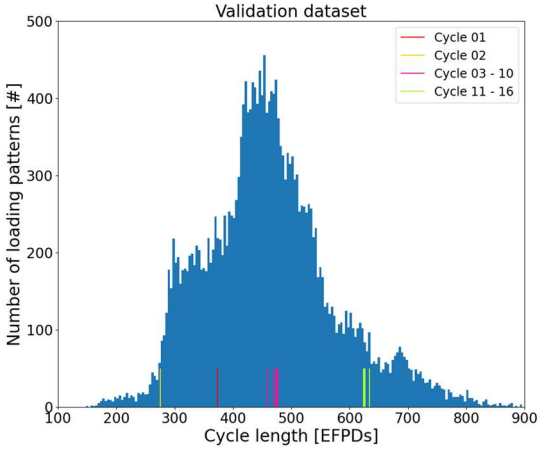

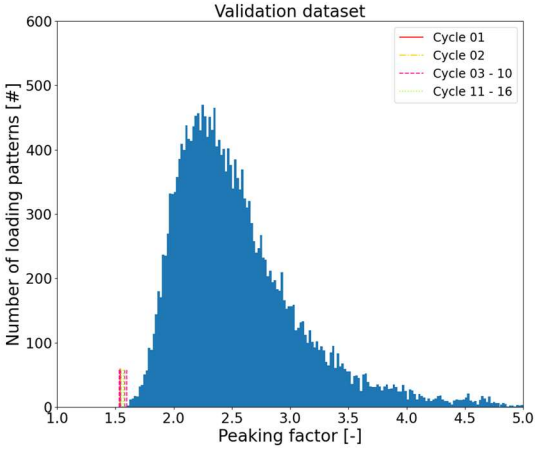

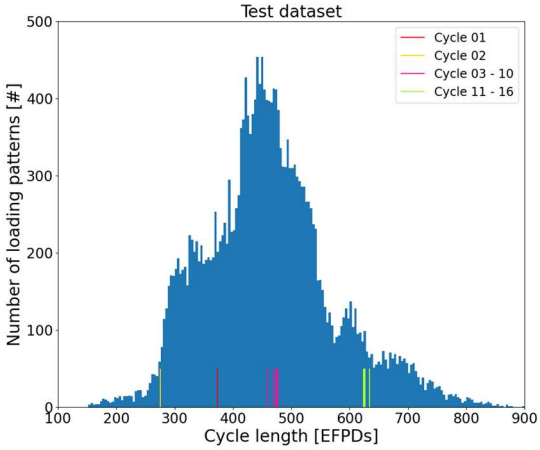

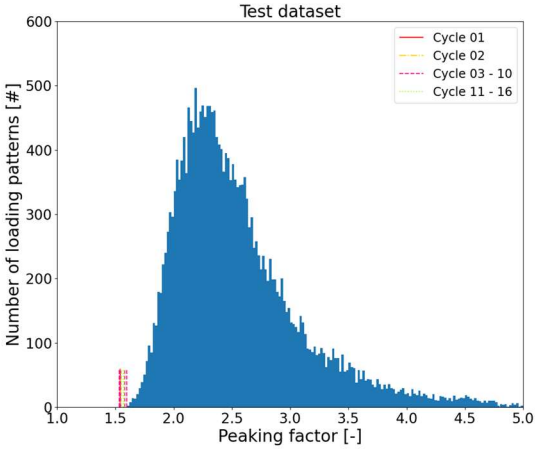

Fig. 4. Distributions of cycle length for the training, validation, and test datasets from the random loading pattern dataset.

Fig. 5. Distributions of peaking factor for the training, validation, and test datasets from the random loading pattern dataset.

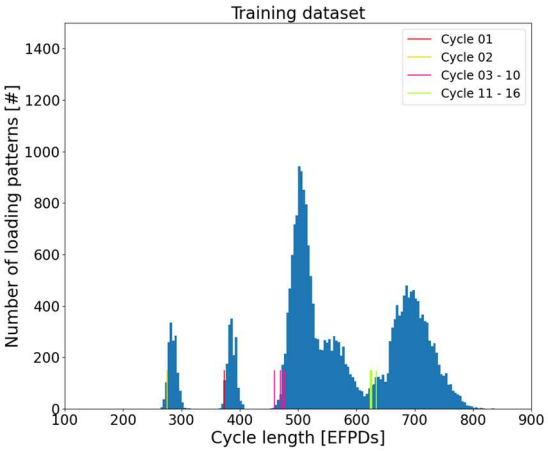

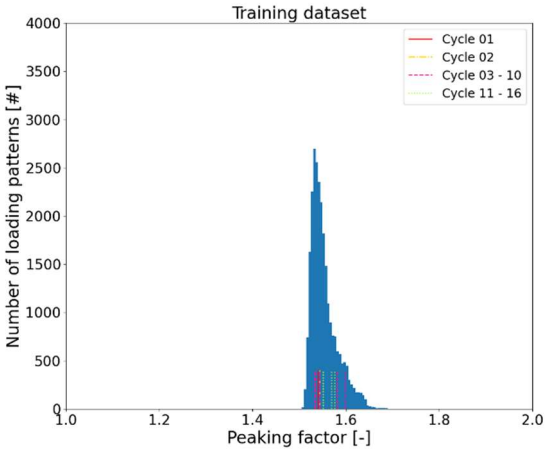

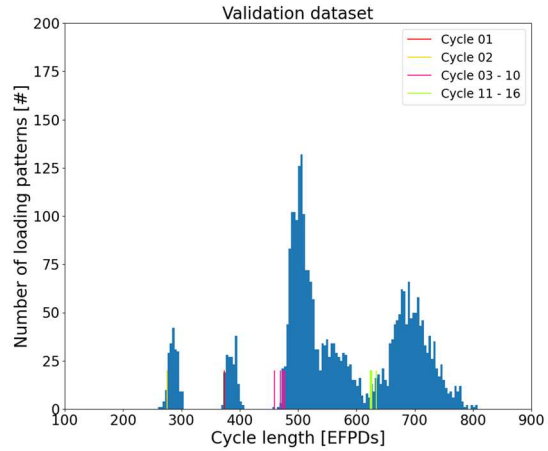

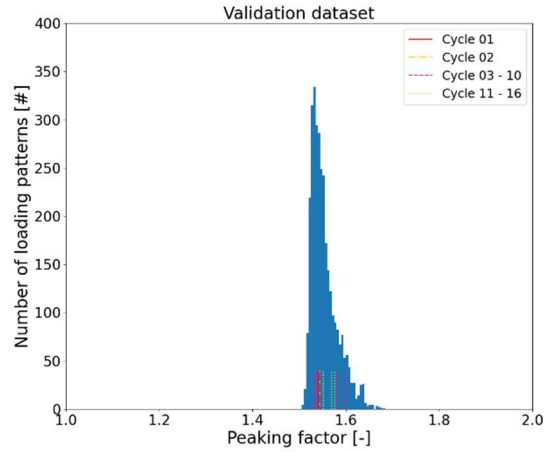

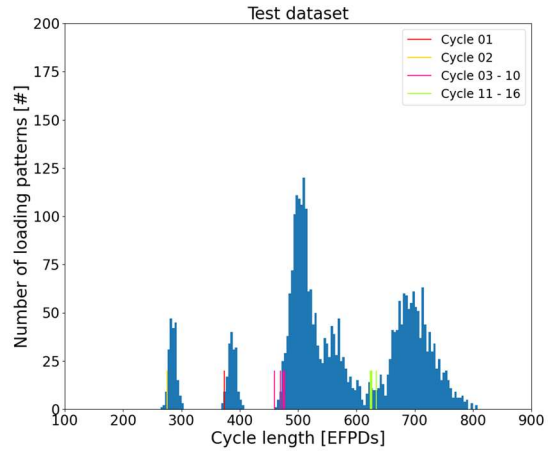

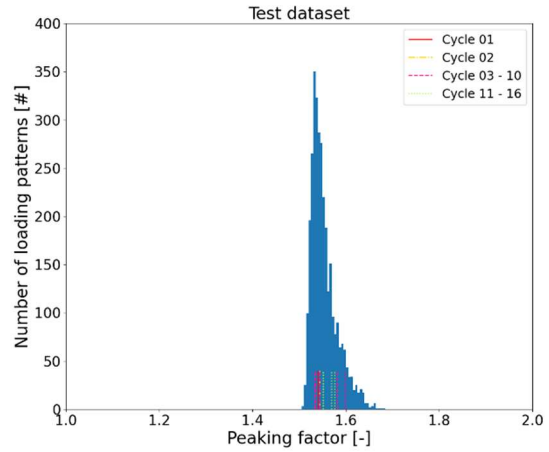

Fig. 6. Distributions of cycle length for the training, validation, and test datasets from the near-optimal loading pattern dataset.

Fig. 7. Distributions of peaking factor for the training, validation, and test datasets from the near-optimal loading pattern dataset.

2.3. Model performance

2.3.1. Model trained with random LP dataset

Figure 8 compares the core performance parameters predicted by the trained model with the corresponding

ST/RK-calculated values using the random LP dataset. The left graph presents the calculated vs. predicted cycle length, while the right graph shows the calculated vs. predicted peaking factor. The red line represents the ideal case where the predicted values perfectly match the ST/RK-calculated values. The black lines indicate the ±2% relative error boundary. In both graphs, the data points are closely aligned with the diagonal reference line, indicating a strong correlation between the model's predictions and the ST/RK results. However, the peaking factor predictions show slightly greater deviation from the reference line than the cycle length predictions, suggesting a relatively higher prediction error for the peaking factor. This is likely due to the higher sensitivity of peaking factor to localized pin-wise heterogeneity, which poses a greater challenge for the model to capture precisely.

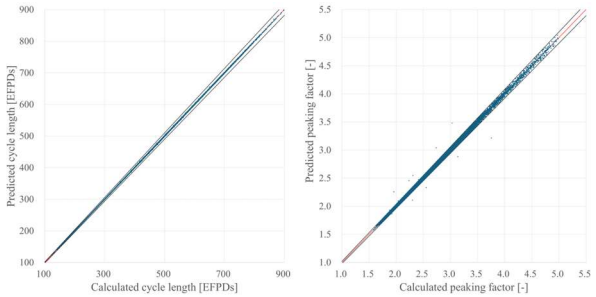

(A) Cycle length evaluation model (B) Peaking factor evaluation model

Fig. 8. Comparison of the core performance parameters between calculated value from ST/RK and predicted value by models trained with random loading pattern dataset.

Table 3 quantifies the relative error between the ST/RK calculations and model predictions for cycle length and peaking factor. The RMS error is 0.12% for cycle length and 0.82% for the peaking factor, confirming that the model achieves high prediction accuracy. The maximum absolute error is 0.85% for cycle length and 19.8% for the peaking factor, further highlighting the increased difficulty in accurately predicting the peaking factor compared to the cycle length. Additionally, 99.7% of the cycle length predictions and 99.9% of the peaking factor predictions fall within three sigma range (3σ) of the ST/RK-calculated values, demonstrating the model's reliability.

Table III: Relative errors of cycle length and peaking factor between ST/RK calculations and predictions of models trained using random loading pattern dataset (All error metrics are expressed in [%]).

Error metric	Cycle length	Peaking factor
RMS error (σ)	0.12	0.82
Maximum error	0.85	19.8
Percentage of errors within $\pm 3\sigma$	99.7	99.9

2.3.2. Fine-tuning with near-optimal LP dataset

To further improve the model's performance in the near-optimal region, we fine-tuned it using the near-optimal LP data generated in the previous step.

Figure 9 compares ST/RK calculation values with the model-predicted values for the test sets of both the random LP dataset and the near-optimal LP dataset before fine-tuning. The red reference line represents perfect prediction accuracy, while the black boundary lines indicate a ±2% relative error range. For cycle length, the model accurately predicts values even in the near-optimal LP dataset. However, for the peaking factor, prediction accuracy significantly decreases in the near-optimal region. This is primarily because the model, initially trained on randomly generated LPs, lacked sufficient data in the low peaking factor region. As a result, it struggled to make accurate predictions for LPs within the near-optimal space, where precise evaluation is critical. This highlights the importance of exposing the model to sufficient high-quality data in the optimal performance region during training.

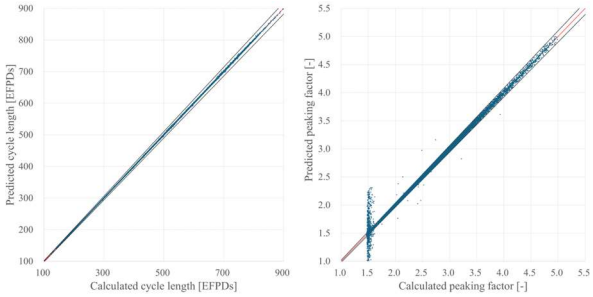

(A) Cycle length evaluation model (B) Peaking factor evaluation model

Fig. 9. Comparison of the core performance parameters between calculated value from ST/RK and predicted value by models before fine-tuning with the near-optimal loading pattern dataset.

Figure 10 presents a similar comparison after fine-tuning the model. The results show that fine-tuning greatly improves the model's accuracy in predicting the peaking factor for near-optimal LPs. Most LP predictions closely align with the red reference line, and the majority fall within the black boundary lines, confirming the enhanced performance of the fine-tuned model.

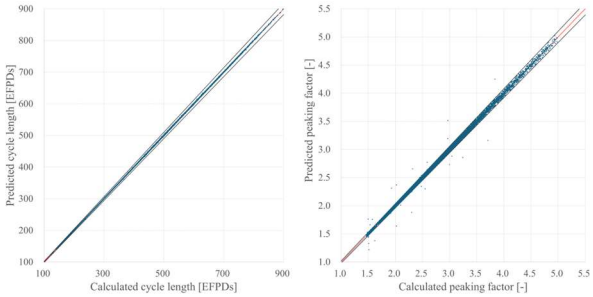

(A) Cycle length evaluation model (B) Peaking factor evaluation model

Fig. 10. Comparison of the core performance parameters between calculated value from ST/RK and predicted value by

models after fine-tuning with the near-optimal loading pattern dataset.

Table 4 quantifies the relative error between ST/RK calculation results and model predictions. The RMS error is 0.12% for cycle length and 0.79% for peaking factor, showing the model achieves high prediction accuracy. The maximum absolute error is 0.86% for cycle length and 19.9% for peaking factors, indicating that some peaking factor predictions still have significant deviations. The ratio of absolute errors within 3σ remains high, at 99.8% for cycle length and 99.9% for peaking factor, confirming the model's reliability. These results demonstrate that the model can accurately infer LPs in the near-optimal region as well.

Table IV: Relative errors of cycle length and peaking factor between ST/RK calculations and predictions of fine-tuned models using the near-optimal loading pattern dataset (All error metrics are expressed in [%]).

Error metric	Cycle length	Peaking factor
RMS error (σ)	0.12	0.79
Maximum error	0.86	19.9
Percentage of errors within $\pm 3\sigma$	99.8	99.9

3. Loading Pattern Optimization

In previous study, random and near-optimal LP generation methods were used to produce data and train the LP evaluation model to assess multi-cycle LPs. Using this model, we performed SA with screening technique.

3.1 Simulated Annealing with a Screening Technique

Figure 11 illustrates the LP search flowchart of the SA algorithm. In each iteration, an LP called X_{cur} is modified by selecting a random assembly and applying one of three modifications: changing the fuel assembly type, swapping positions with another assembly, or rotating the assembly one to three times by 90°. When changing assembly types or swapping positions, octant symmetry is maintained. The LP evaluation model then predicts the cycle length and peaking factor of the newly generated LP, X_{new}. Based on these values, the objective function J_{new} is calculated using equation 1.

$$(1)\ J(X) = J_{PF}(X) + J_{CYC}(X)$$

In the equations, PF denotes the predicted peaking factor and CYC denotes the predicted cycle length, both obtained from the ViT-based evaluation models. $PF(X)$ and $CYC(X)$ represent the values predicted for the candidate loading pattern X, while PF_{ref} and CYC_{ref} are the corresponding values of the reference loading pattern.

$$(2)\ J_{PF}(X) = \frac{1}{PF_{ref}} \frac{\left(PF(X) - PF_{ref}\right)^3}{\left|PF(X) - PF_{ref}\right|}$$

$$(3)\ J_{CYC}(X) = \frac{1}{CYC_{ref}} \frac{\left(CYC(X) - CYC_{ref}\right)^3}{\left|CYC(X) - CYC_{ref}\right|}$$

The valid range (J_{min}, J_{max}) is determined within a 3σ range, where σ represents the RMS error of the model's prediction across the entire dataset. Equation 4 illustrates how the valid range is defined, with $\overline{\Delta J}$ representing the bias of the model's predictions for the total dataset.

$$(4)\ J_{max/min} = J_{new} + \overline{\Delta J} \pm 3\sigma$$

The acceptance criterion J_{acp} is derived from current value J_{cur}, like equation 5.

$$(5)\ J_{acp} = J_{cur} - C \ln \xi$$

If J_{min} is greater than J_{acp}, X_{new} is rejected, and X_{cur} is used to generate a new LP. If J_{max} is lower than J_{acp}, which means this LP is sufficiently good, the adjusted objective function, $\tilde{J}_{new} = J_{new} + \overline{\Delta J}$, is evaluated. After that, X_{cur} is updated to X_{new} and J_{cur} is set by $\tilde{J}_{new}$. If J_{acp} falls within the range of J_{min} to J_{max}, a 3D deterministic calculation is performed to obtain the exact J_{new}, which is then used to update the model's mean prediction error and sigma. If J_{new} is lower than J_{cur}, X_{cur} is then updated to X_{new}, and the process repeats. If the repeat number exceeds 200 times or X_{cur} updates 10 times, the temperature C is reduced, and the process restarts. If no X_{cur} updates occur after 200 times of repeat, the SA process terminates.

Table 5 presents the SA results for Cycle 1 LP optimization using an AI-based screening technique. The efficiency represents the percentage of LP evaluations performed by AI. ST refers to the Screening Technique. The results show that, on average, 99.9% of LP evaluations were overseen by the AI model, significantly reducing computational time. Each SA process involved approximately 10,000 LP evaluations, with only about 10 evaluations performed using RAST-K, demonstrating the high efficiency of AI-driven screening.

Notably, despite being conducted on a single core without parallel processing, each SA process required only an average of 1.4 hours. In contrast, without AI-based screening, the estimated computation time per process would have ranged from 800 to 900 hours. This substantial reduction in computational cost underscores the effectiveness of AI-based LP evaluation in accelerating SA optimization, making multi-cycle LP optimization practically feasible within a drastically shorter time.

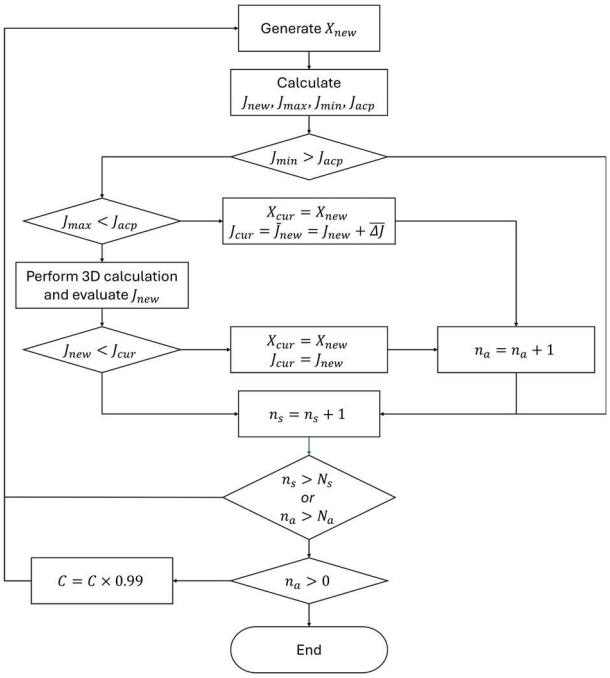

Fig. 11. Flowchart of the optimal loading pattern search using simulated annealing with the screening technique [3].

Table V: Simulated Annealing results using AI-based Screening Technique for Cycle 1 LP optimization.

Index	No. of LP evaluation			Efficiency* [%]	Calculation time [hr.]	Estimated time w/o screening technique [hr.]
	RAST-K	AI	Total			
1	11	10138	10149	99.9	1.2	845.8
2	12	10203	10215	99.9	1.3	851.3
3	6	9683	9689	99.9	0.8	807.4
4	9	9958	9967	99.9	1.0	830.6
5	11	10305	10316	99.9	1.2	859.7
6	14	10504	10518	99.9	1.5	876.5
7	17	10853	10870	99.8	1.7	905.8
8	10	10017	10027	99.9	1.1	835.6
9	19	11214	11233	99.8	1.9	936.1
10	19	10914	10933	99.8	1.9	911.1
11	9	9934	9943	99.9	1.0	828.6
12	10	10184	10194	99.9	1.1	849.5
13	5	9535	9540	99.9	0.7	795.0
14	19	11111	11130	99.8	1.9	927.5
15	9	10048	10057	99.9	1.0	838.1
16	18	11079	11097	99.8	1.8	924.8
17	18	10858	10876	99.8	1.8	906.3
18	20	11011	11031	99.8	2.0	919.3
19	17	10938	10955	99.8	1.7	912.9
20	5	9571	9576	99.9	0.7	798.0
Avg.	12.9	10402.9	10415.8	99.9	1.4	868.0

4. Results

In Cycle 1, a new LP for the SA method is generated using fresh fuel assemblies of A0 to D2. Fresh fuel assemblies for Cycle 2 are E0 to E2. LP of Cycle 1 has a cycle length of 373.2 EFPDs and a peaking factor of 1.5412 and LP of Cycle 2 has a cycle length of 275.8 EFPDs and a peaking factor of 1.5442. Figure 12 shows the NDR fuel LP of Cycles 1 and 2. In the following figures, blue indicates fresh fuel and green shows the assembly is once burnt.

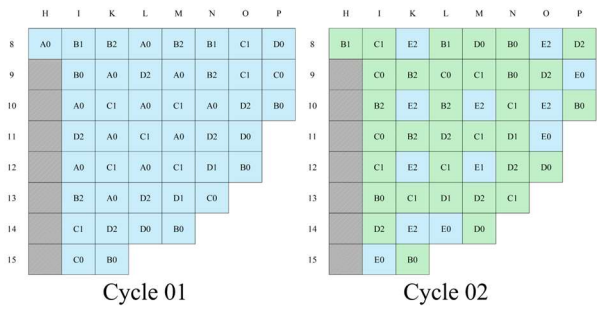

Fig. 12. Reference fuel loading patterns of the Cycle 1 and Cycle 2 cores of Hanbit Unit 3 from the Nuclear Design Report [6].

There are three methods for generating new LPs in SA. The first method preserves the number of fuel assembly types, allowing only swapping and rotation of assemblies. The second method preserves only the enrichment value while allowing the assembly type change only their BP fractions are identical. For example, B0 and B2 have the same fuel enrichment but differ in the number of BP rods. In this case, B0 can be replaced with B2, and vice versa, within the LP. The third method allows all the changes for fresh fuel assembly within the fresh fuel assemblies used in the current cycle. However, the last approach alters the total U-235 mass in the core, potentially affecting fuel costs.

In Cycle 2, the burnt fuel assemblies from the previous cycle may not exist in the optimized LP for Cycle 1, and thus they are placed to match their previous positions and rotation numbers. the burnup values of the burnt fuel assembly are set based on the results of the Cycle 1 optimal LP using the same LP shuffling method. The red boxes in the following figures highlight how the placement of fresh fuel assemblies has changed.

4.1 Optimal LP with no assembly type change

In this approach, the number of fuel assembly types remains unchanged, allowing only swapping and rotation of assemblies within the LP. By preserving the original assembly types from the NDR LP, this method optimizes the loading pattern while maintaining the initial fuel composition. Figure 13 represents optimized LPs.

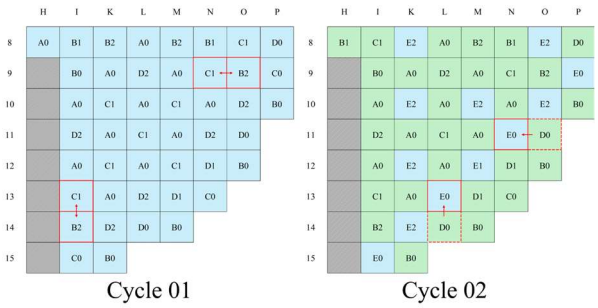

Fig. 13. Optimal fuel LP while maintaining the number of fuel assembly types for Cycle 1 and Cycle 2 core of Hanbit Unit 3. Red boxes indicate changed positions of fresh fuel assemblies after optimization.

In Cycle 1, the optimized LP achieved a cycle length of 382.9 EFPDs and a peaking factor of 1.4750, showing improvement over the NDR core parameters (373.2 EFPDs, 1.5412). In Cycle 2, the optimized LP achieved a cycle length of 277.2 EFPDs and a peaking factor of 1.5356. This is an improvement over the NDR core parameters (275.8 EFPDs, 1.5442), showing a slight increase in cycle length and minor reduction in a peaking factor. This approach resulted in longer cycle lengths and reduced peaking factors. The improvement in a peaking factor indicates better power distribution, potentially enhancing core safety margins.

4.2 Optimal LP with BP fraction change

In this approach, the enrichment of each fuel assembly remains unchanged, while modifications to the BP mass fraction are allowed. This method provides greater flexibility in power distribution optimization without altering the overall U-235 mass in the core. Figure 14 illustrates the optimal LP generated by SA.

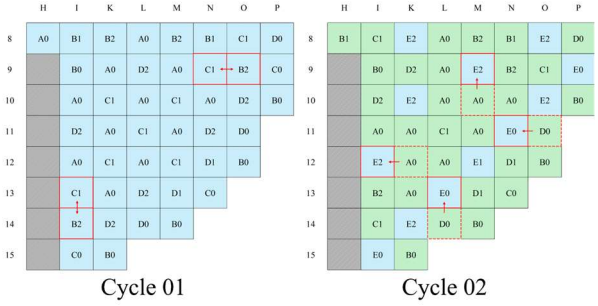

Fig. 14. Optimal fuel LP while BP fraction is changed by type conversion for Cycle 1 and Cycle 2 core of Hanbit Unit 3. Red boxes highlight assemblies with modified burnable poison fractions.

For Cycle 1, the SA-optimized LP with BP fraction changes was identical to the LP obtained using the previous method, resulting in a cycle length of 382.9 EFPDs and a peaking factor of 1.4750. This suggests that as long as the enrichment remains unchanged, BP modifications have minimal impact on optimization in Cycle 1.

For Cycle 2, the BP fraction change approach produced an optimized LP with a cycle length of 277.4 EFPDs and a peaking factor of 1.5422. This is an improvement over the NDR core parameters (275.8 EFPDs, 1.5442), showing a slight increase in cycle length and minor reduction in a peaking factor. Although the enhancement is less pronounced than in Cycle 1, the method still demonstrates its potential for optimizing power distribution while maintaining core reactivity.

4.3 Optimal LP with enrichment and BP fraction change

This approach allows modifications to both the fuel enrichment and the BP mass fraction of fuel assemblies, providing the highest flexibility in optimizing core performance. Unlike the previous methods, which either preserved assembly types or only adjusted BP fractions, this approach enables more extensive reconfiguration of the LP by selecting assemblies with different U-235 enrichment levels and BP content. However, altering enrichment directly impacts the total U-235 mass in the core, potentially affecting fuel costs and reactivity control. Figure 15 illustrates the optimal LP generated by SA.

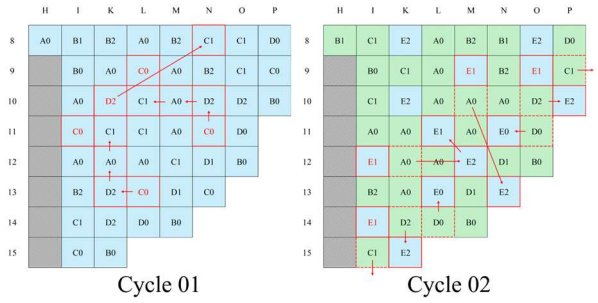

Fig. 15. Optimal fuel LP while enrichment and BP fraction are changed by type conversion for Cycle 1 and Cycle 2 core of Hanbit Unit 3. Red boxes show assemblies with both enrichment and BP fraction changes.

For Cycle 1, the SA-optimized LP with both enrichment and BP fraction changes resulted in a cycle length of 389.5 EFPDs and a peaking factor of 1.5373. Compared to the NDR core parameters (373.2 EFPDs, 1.5412), this approach significantly extended the cycle length while maintaining a comparable peaking factor.

For Cycle 2, the optimized LP achieved a cycle length of 277.3 EFPDs and a peaking factor of 1.5439, compared to the NDR core values (275.8 EFPDs, 1.5442). While the improvement in cycle length is modest, the approach provides enhanced flexibility in optimizing reactivity distribution.

5. Conclusions

In this study, a ViT-based multi-cycle prediction model was developed and validated for evaluating the fuel LP of the OPR-1000 reactor. Compared to conventional three-dimensional deterministic methods,

the proposed model significantly reduced computational costs while maintaining high prediction accuracy and demonstrating its effectiveness in multi-cycle LP evaluations.

The model achieved strong agreement with ST/RK calculations, with an RMS error of 0.12% for cycle length and 0.82% for peaking factor when trained with randomly generated data. However, due to the limited representation of low peaking factor LPs in the random dataset, prediction accuracy was reduced in the near-optimal region. To overcome this, we generated and incorporated near-optimal LP data through SA, which enabled fine-tuning of the model. As a result, the fine-tuned model achieved improved accuracy, especially in the critical low peaking factor region, with an RMS error of 0.79%.

Using the fine-tuned ViT evaluation model, LPs for Cycle 1 and Cycle 2 were successfully optimized through SA combined with a screening technique. This AI-based screening approach enabled over 99.9% of LP evaluations to be conducted by the model instead of full 3D calculations, reducing the optimization time from hundreds of hours to just 1–2 hours on a single core without parallelization.

All three LP generation methods assessed in this study led to extended cycle lengths and reduced peaking factors compared to the NDR LPs, confirming the effectiveness of the proposed methodology. The integration of a pin-wise ViT evaluation model with SA not only enhanced fuel cycle performance and core safety margins but also demonstrated significant potential for accelerating and automating the LP optimization process.

Future research will focus on identifying optimal LPs beyond Cycle 3 and developing predictive models capable of estimating the best LP configurations for subsequent cycles. In parallel, efforts are ongoing to apply AI throughout the entire optimization process by developing an optimal LP generative model, with the aim of achieving full automation and further reducing computational costs in LP optimization.

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 Human Resources Program in Energy Technology of the Korea Institute of Energy Technology Evaluation and Planning (KETEP), which was funded by the Ministry of Trade, Industry & Energy (MOTIE of Korea) (No. RS-2024-00398425).

REFERENCES

[1] P. W. Poon and G. T. Parks, Application of Genetic Algorithms to In-Core Nuclear Fuel Management Optimization, Proc. Topl. Mtg. Mathematical Methods and Supercomputing in Nuclear Applications, Vol. 2, p. 777, 1993.
[2] D. J. Kropaczek and P. J. Turinsky, In-Core Nuclear Fuel Management Optimization for Pressurized Water Reactor Using Simulated Annealing, Nucl. Technol., vol. 95, p. 9-32, 1991.
[3] 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, vol. 162, p. 134-147, 2009.
[4] H. Jang, H. C. Shin, D. Kim and H. C. Lee, Application of Convolutional Neural Network to Fuel Loading Pattern Optimization by Simulated Annealing, Korean Nuclear Society Virtual Autumn Meeting, 2020.
[5] S. Jeong and H. C. Lee, Loading Pattern Optimization for OPR-1000 Using Simulated Annealing with Pin-wise Vision Transformer Based Screening Technique, Transactions of the Korean Nuclear Society Autumn Meeting, 2024.
[6] Nuclear Design Report for Yonggwang Nuclear Power Plant Unit 3 Cycle 2, KAERI/TR-605/96, Korea Atomic Energy Research Institute, 1996.
[7] J. Lee and H. C. Lee, Loading pattern design and economic evaluation for 24-month cycle operation of OPR-1000 in Korea, Vol. 55(3), p. 1167-1180, Nuclear Engineering and Technology, 2022.
[8] J. Park, J. Jang, H. Kim, Et Al., RAST-K v2- Three Dimensional Nodal Diffusion Code for Pressurized Water Reactor Core Analysis, Energies, 13: 6324, 2020.