Development of a Unified Loading Pattern Evaluation Model Using a Vision Transformer

Seongjin Jeong and Hyun Chul Lee*
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, feature-wise linear modulation, cross vision transformer, simulated annealing

1. Introduction

Loading pattern (LP) optimization extends the cycle length of a PWR while satisfying design limits such as the power peaking factor. Because the search space is combinatorial and each candidate requires a three-dimensional (3D) core calculation, metaheuristic methods such as simulated annealing (SA) [1] demand an impractical number of evaluations. Screening-based SA relieves this cost by predicting the objective value with a fast surrogate and invoking the 3D calculation only for candidates whose acceptance decision the prediction cannot determine [2]. In our previous work, a pin-wise vision transformer (pin-wise ViT) surrogate was shown to screen about 99.9% of the candidates and to complete a single-cycle optimization of the Hanbit Unit 3 (YGN3) cycle 1 core within a few hours [3].

However, there remain two limitations. First, the surrogate consisted of two single-task models predicting only the cycle length and the maximum F_{xy} in a cycle (max-F_{xy}), so that the remaining core parameters needed for a design-level assessment still required a separate 3D calculation. Second, only a single cycle was optimized, whereas a realistic design must maintain inter-cycle consistency because the burnt fuel assemblies (FAs) of a cycle are inherited from the optimized LP of the preceding cycle.

This paper presents a unified vision transformer (unified ViT) that predicts, from a single LP image, both the scalar quantities used by the SA objective function and the burnup-dependent parameters used for LP evaluation. The model is employed as the screening surrogate of the multi-objective adaptive SA [2, 4], in which the reference core parameters are calculated by STREAM/RAST-K v2 [5], to optimize ten consecutive OPR-1000 cycles sequentially.

2. Methodology

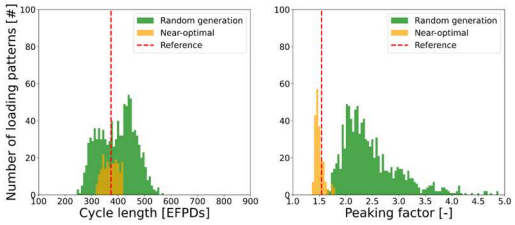
2.1 Loading pattern datasets

Random generation (RG) and near-optimal (NO) datasets were constructed for ten consecutive cycles. Cycle 1 and cycle 2 follow the Hanbit Unit 3 (YGN3) nuclear design reports [6], and cycles 3 to 10 are set as the cores of an LP design study for 18-month equilibrium cores [7]. The FA types used in the nuclear design reports (NDR) and the LP design study are listed in Table I.

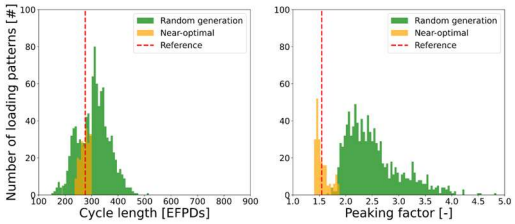
Table I: Specifications of fuel assemblies in the reference loading patterns [6, 7].

FA type	Fuel enrichment [wt.%^{235}U]		No. of Rods		Burnable poison fraction [wt.%Gd_2O_3]
	Normal	Zoned	Zoned	BP	
A0	1.30	-	-	-	-
B0	2.37	-	-	-	-
B1	2.36	1.30	52	8	4.0
B2	2.37	-	-	4	4.0
C0	2.87	2.35	52	-	-
C1	2.87	2.36	52	8	4.0
D0	3.35	2.87	52	-	-
D1	3.36	2.85	52	8	4.0
D2	3.35	2.87	100	8	4.0
E0	4.08	3.61	52	-	-
E1	4.08	3.61	52	8	6.0
E2	5.60	3.11	52	8	6.0
FC	2.20	-	-	-	-
F0	4.65	4.10	52	-	-
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

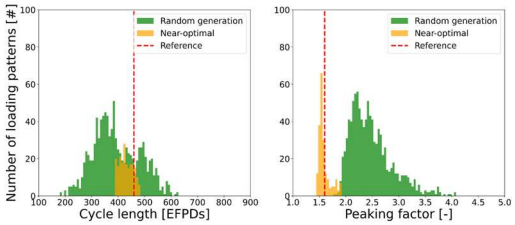
For each cycle, 10,000 RG cases were generated under a relaxed constraint (max-F_{xy} < 5.0), using a randomly generated LP of the preceding cycle for the burnt FAs with a fixed number of fresh FAs, and 2,000 NO cases were collected from SA runs screened by a unified ViT pre-trained on the RG dataset only. The totals are 100,000 RG and 20,000 NO cases, split into training and test subsets at a ratio of 9:1. Figure 1 shows the cycle-wise distributions of the cycle length and the peaking factor for the RG and NO datasets. In the peaking factor distributions, the RG dataset peaks at around 2.0 and contains almost no LPs in the low-peaking-factor region near the reference core value, indicated by the red dashed line in each plot, where the NO dataset is concentrated. Because the RG-pretrained model showed a markedly larger error there, the NO dataset was used for fine-tuning.



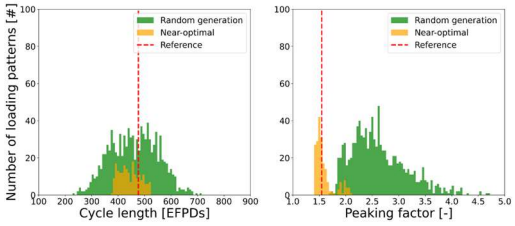
(a) Cycle 01



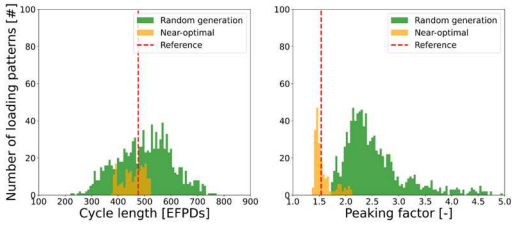
(b) Cycle 02



(c) Cycle 03



(d) Cycle 05



(e) Cycle 10

Fig. 1. Cycle-wise distributions of cycle length (left) and peaking factor (right) in the dataset for representative cycles.

2.2 Unified Vision Transformer model

An overview of the unified ViT is depicted in Figure 2. The model takes a quarter-core image at pin resolution with three channels — the pin-wise enrichment, the pin-wise burnable poison (BP) fraction, and the box-wise burnup [3] — and differs from the previous single-task ViT models in three respects.

The first is dual-scale tokenization. Patches of 4×4 and 8×8 pin size are embedded in parallel and fused by cross-attention [8], replacing the single 8×8 patch embedding. The fine-scale tokens preserve the intra-assembly heterogeneity that governs the peaking factor, while the coarse-scale tokens describe the global power tilt.

The second is burnup conditioning by feature-wise linear modulation (FiLM) [9]. The burnup channel modulates the enrichment and BP-fraction features through FiLM rather than entering as a raw input channel, reflecting the physical role of burnup as a modifier of the local reactivity of a given FA composition.

The third is a unified multi-task head. A single model predicts the scalar quantities used by the SA objective function (cycle length, max-Fxy) and the burnup-dependent parameters — the critical boron concentration (CBC), the axial shape index (ASI), five peaking factors (Fq, Fr, FdH, Fxy, Fz), and the radial box-wise power map, tabulated as functions of burnup — at once, replacing the previous two single-task models.

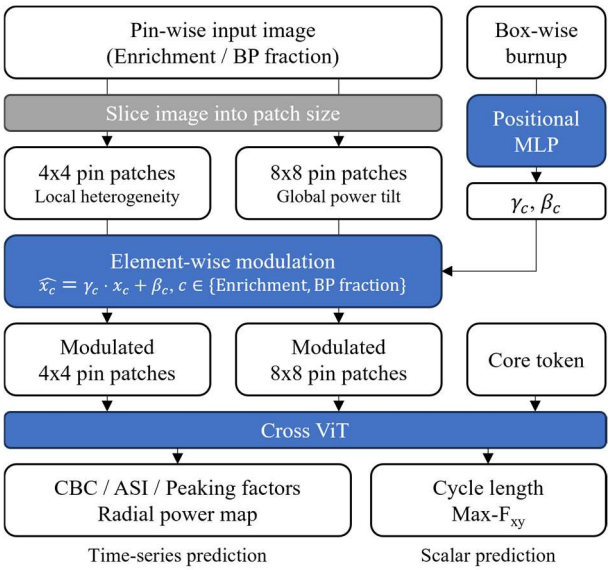


Fig. 2. Overview of the unified ViT for loading pattern evaluation.

The model was trained with an MSE loss using AdamW and a ReduceLROnPlateau scheduler, pre-trained on the RG dataset and fine-tuned on the NO dataset to improve accuracy in the near-optimal region where the search converges. Training took about 20 h on a single NVIDIA RTX 5060 Ti (16 GB) GPU.

2.3 Sequential multi-cycle optimization

Candidates were generated with the type-fix option, which preserves the number of each fresh FA type in the LP so that the total ^{235}U mass is unchanged and only positional shuffling is searched. The SA objective function follows the multi-objective formulation with the discontinuous penalty function [2] and is defined as

$$(1)\ J(x) = J_{CL}(x) + J_{PF}(x)$$

$$(2)\ J_{CL}(x) = \begin{cases} 1 + \dfrac{(CL(x) - CL_{\text{lim}})^2}{\overline{\delta_{CL}^2}} & (CL(x) < CL_{\text{lim}}) \\ 0 & (CL(x) \geq CL_{\text{lim}}) \end{cases}$$

$$(3)\quad J_{PF}(x)=\begin{cases}1+\dfrac{(PF(x)-PF_{\lim})^2}{\overline{\delta_{PF}^2}} & (PF(x)>PF_{\lim})\\ 0 & (PF(x)\le PF_{\lim})\end{cases}$$

where x : Loading pattern,
$CL(x)$: Cycle length,
$PF(x)$: Peaking factor,
$CL_{\lim}$: Target cycle length,
$PF_{\lim}$: Peaking factor limit and
$\overline{\delta_{CL}^2}$, $\overline{\delta_{PF}^2}$: Normalization factor

The normalization factor is the mean squared deviation of each parameter from its target or limit over the test subset of the corresponding cycle. Each term is zero when the target or limit is satisfied and positive otherwise, so that minimizing J(x) maximizes the cycle length within the peaking factor limit. The ten cycles were optimized sequentially as shown in Figure 3: the initial LP of each cycle was a randomly generated LP based on the reference, the burnt FAs were taken from the optimized LP of the preceding cycle, and the LP with the longest cycle length satisfying the peaking factor limit during an SA process was selected as the optimal LP and passed to the next cycle. The target cycle length and the peaking factor limit were taken from the NDR cores for cycle 1 and cycle 2. For cycles 3 to 10, the peaking factor limit was set to 1.60 for cycle 3 and cycle 4 and 1.55 for cycles 5 to 10, and the target cycle length was set to 490 EFPDs, which is intentionally higher than the reference cycle lengths (459.8–476.4 EFPDs) so that the cycle length term remains active throughout the search and drives the SA toward the longest cycle length attainable within the peaking factor limit.

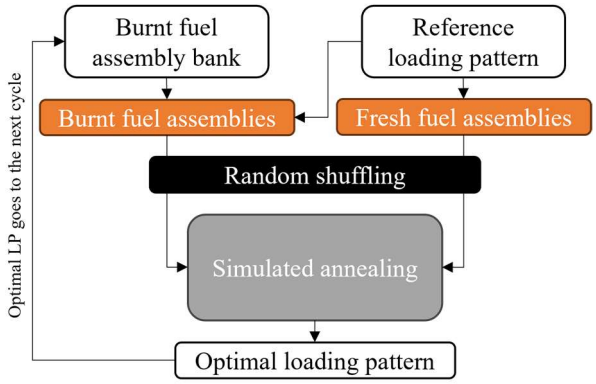


Fig. 3. Overview of the sequential multi-cycle optimization using simulated annealing.

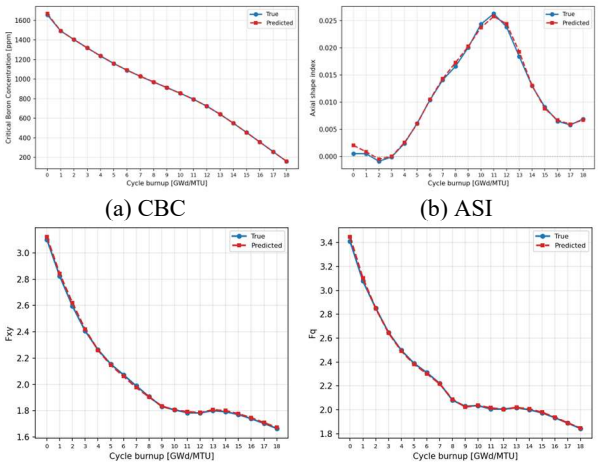
3. Results

3.1 Prediction accuracy

Table II compares the RMS relative prediction error against the RAST-K calculation for the test subset, and Figure 4 shows the predicted CBC, ASI, F_{xy}, and F_q of a random LP compared with the RAST-K calculation in the test subset. The unified ViT reduced the cycle length error from 0.16% to 0.04% and the max-Fxy error from 0.86% to 0.41% relative to the previous single-task models and additionally predicted the burnup-dependent parameters with RMS errors below 0.7%. The improved max-F_{xy} precision is directly relevant to the screening, since the width of the ±3σ band determines how many candidates are passed to the 3D calculation.

Table II: RMS relative errors of pin-wise and unified ViT predictions with respect to RAST-K calculations for the test subset.

Metric		Pin-wise ViT [%]	Unified ViT [%]
Scalar prediction	Cycle length	0.16	0.04
	Max-F_{xy}	0.86	0.41
Time-series prediction	CBC	-	0.13
	ASI	-	0.20
	Peaking factors	-	0.43
	Radial power map	-	0.69



(a) CBC (b) ASI

(c) F_{xy} (d) F_q

Fig. 4. Comparison of model predictions and RAST-K calculations for a random loading pattern in the test subset.

Table III summarizes the average number of evaluations, the screening efficiency, and the process time for each cycle. Here, the process time is the wall-clock time of an entire SA run, which includes both the unified ViT inference on the GPU and the RAST-K calculations on the CPU. Of the 34,000–58,000 candidates evaluated per run, 99.8–99.9% were decided by the surrogate alone, leaving only 24–79 3D core calculations, so that a single run required 1.2–8.1 h against an estimated 1,091–5,305 h with RAST-K alone. The number of 3D calculations grows with the cycle index, reflecting the larger number of burnt FA types and the tighter peaking factor limit of the later cycles.

Table III: Average values of 10 simulated annealing runs using the unified ViT for cycles 1–10.

Cycle	RAST-K calls	Unified ViT calls	Efficiency [%]	Process time [h]	Est. process time w/o screening [h]
1	28	34,432	99.9	1.6	1,091
2	24	40,110	99.9	1.2	1,167
3	33	38,755	99.9	3.4	3,017
4	58	40,085	99.9	5.8	3,420
5	65	40,138	99.8	6.7	3,648
6	73	41,585	99.8	7.5	3,780
7	79	50,111	99.8	8.1	4,554
8	59	40,971	99.9	6.1	3,723
9	62	58,405	99.9	6.7	5,305
10	67	56,228	99.9	7.2	5,108

Figure 5 shows the process time, cycle length, and peaking factor of the optimal LPs of the ten cycles, and Table IV lists representative LPs with the reference values. In all cycles except cycle 4, the optimal LP achieved both a longer cycle length and a lower peaking factor.

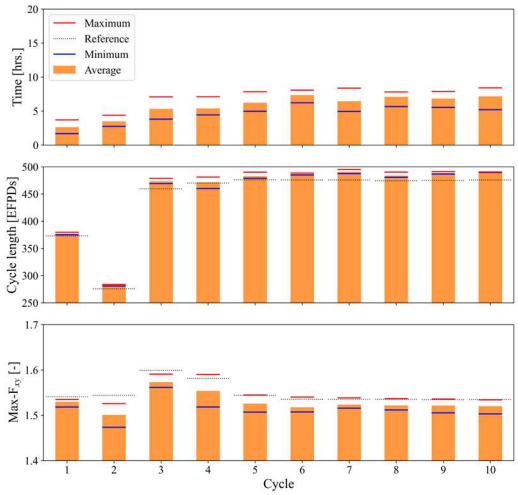
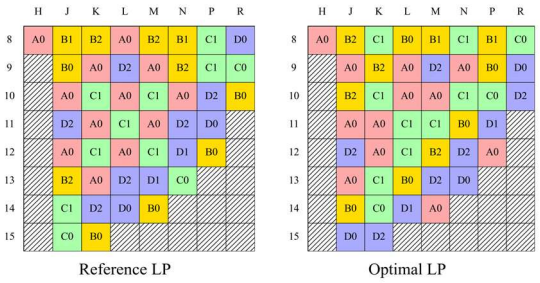


Fig. 5. Cycle-wise simulated annealing process time (top) and core performance metrics of the optimal loading patterns: cycle length (middle) and peaking factor (bottom).

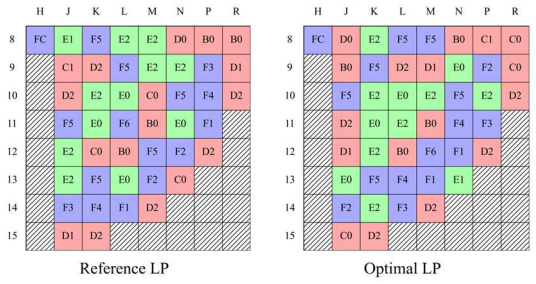
Table IV: Cycle length and peaking factor of cycle-wise reference and representative loading patterns.

Cycle	Cycle length [EFPDs]		Max-F_{xy}	
	Reference	Optimal	Reference	Optimal
1	373.2	376.2	1.5412	1.5317
2	275.8	279.9	1.5442	1.5012
3	459.8	476.7	1.5994	1.5665
4	470.3	460.9	1.5812	1.5599
5	476.4	477.6	1.5443	1.5398
6	476.0	488.9	1.5356	1.5236
7	475.9	489.1	1.5354	1.5339
8	474.3	480.3	1.5355	1.5276
9	475.4	479.7	1.5344	1.5285
10	475.9	490.8	1.5349	1.5063
Sum	4,433.0	4,500.1	-	

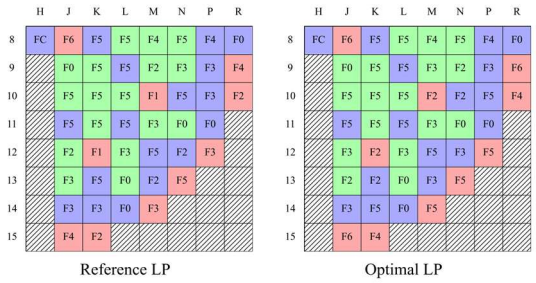
Figure 6 shows the representative LPs that yielded the longest cumulative cycle length among the SA runs. In cycle 1, the A0-type FAs are arranged in a checkerboard-like pattern in the reference LP, but the pattern is slightly broken in the optimized LP. In the subsequent cycles, the twice-burnt FAs are mostly placed at the core periphery, and the fresh FAs are distributed slightly more inward than in the reference LPs, indicating that the optimization consistently reproduces a low-leakage loading strategy while inheriting the burnt FA inventory from the preceding cycle.



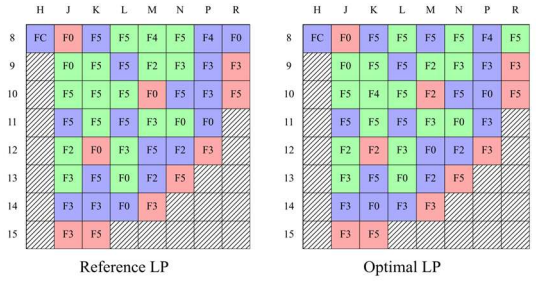
(a) Cycle 01



(b) Cycle 03



(c) Cycle 06



(d) Cycle 10

Fig. 6. Representative loading patterns yielding the longest cumulative cycle length among the simulated annealing runs. Colors denote assembly type for cycle 1 and burnup history for the subsequent cycles (blue: fresh; green: once-burnt; red: twice-burnt).

The cycle length of cycle 4 was shorter than that of its reference LP. This follows from the sequential optimization: the long cycle achieved in cycle 3 consumed more of the reactivity of the FAs inherited by cycle 4, leaving a less reactive burnt FA inventory. Greedy maximization of a single cycle therefore does not necessarily maximize the performance of the following cycle, which motivates a multi-cycle objective.

4. Conclusions

A unified vision transformer combining dual-scale tokenization, FiLM burnup conditioning, and a multi-task head was developed for the LP evaluation of OPR-1000 cores and trained with 100,000 randomly generated and 20,000 near-optimal LPs covering ten cycles. It predicted the cycle length and max-Fxy with RMS relative errors of 0.04% and 0.41%, reduced by factors of approximately four and two compared with the previous single-task models, and additionally predicted the CBC, ASI, peaking factors, and radial power map with errors below 0.7%, so that an LP can be assessed at the design level without a separate 3D calculation.

Used as the screening model of the adaptive SA over ten consecutive cycles, it screened 99.8–99.9% of the candidates without a 3D calculation and reduced the process time to 1.2–8.1 h per run. The optimized LPs generally showed a longer cycle length and a lower peaking factor than the reference LPs while maintaining inter-cycle consistency, although the sequential optimization produced a shorter cycle length than the reference in cycle 4.

Future work will address direct LP generation, in which a transformer encoder–decoder produces a complete LP in a single pass, conditioned on target parameters, under the fuel-budget constraint, and reinforcement learning, in which LP construction is treated as a sequential decision process with the policy trained using rewards evaluated by the unified ViT.

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 Climate, Energy & Environment(MCEE) of the Republic of Korea (No. RS-2024-00398425).

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, Nucl. Sci. Eng., 162(2), 134–147 (2009).

[3] S. Jeong and H. C. Lee, Effect of Surrogate Model on Screening-based Simulated Annealing for OPR-1000 Loading Pattern Optimization, Korean Nuclear Society Spring Meeting (2026).

[4] H. C. Lee, H. J. Shim, and C. H. Kim, Parallel Computing Adaptive Simulated Annealing Scheme for Fuel Assembly Loading Pattern Optimization in PWRs, Nucl. Technol., 135, 39–50 (2001).

[5] 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).

[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, Nucl. Eng. Technol., 55(3), 1167–1180 (2023).

[8] C. Chen, Q. Fan, and R. Panda, CrossViT: Cross-Attention Multi-Scale Vision Transformer for Image Classification, ICCV (2021).

[9] E. Perez, F. Strub, H. de Vries, V. Dumoulin, and A. Courville, FiLM: Visual Reasoning with a General Conditioning Layer, AAAI (2018).