Optimization of a Molten Salt Reactor Core Design Using MLP Surrogate Model and Genetic Algorithm

Chanhwi An[a] , Yoonpyo Lee [a,b*]
[a]Department of Nuclear Engineering, Hanyang University, 222, Wangsimni-ro, Seongdong-gu, Seoul, 04763, Republic of Korea.
[b]The Grainger College of Engineering, Nuclear, Plasma & Radiological Engineering, University of Illinois Urbana-Champaign, Urbana, IL, USA
[*]*Corresponding author: yoonpyo2@illinois.edu*

***Keywords :** Molten Salt Reactor, Multi-Layer Perceptron, Genetic Algorithm, Reactor Optimization, Nuclear Safety*

1. Introduction

As the global imperative for carbon neutrality intensifies, nuclear energy has emerged as a critical baseload power source capable of ensuring energy security while mitigating climate change [1]. In particular, the Generation IV International Forum has identified the Molten Salt Reactor (MSR) as a promising candidate due to its inherent safety features [2], such as the negative temperature coefficient of reactivity and the ability to operate at low pressures.

Among the MSR concepts, the Molten Salt Fast Reactor (MSFR) has garnered significant attention for its potential to achieve a closed fuel cycle and high fuel sustainability through a superior neutron economy [3]. However, the neutronic design and analysis of MSFRs present unique challenges compared to conventional nuclear reactors. The liquid fuel capability requires a precise evaluation of the interplay between geometric configurations and fuel composition, often necessitating high-fidelity transport codes like Monte Carlo simulations [4]. While Monte Carlo codes, such as OpenMC or MCNP [5,6], provide accurate results by simulating individual particle histories, they are computationally expensive.

To mitigate this burden, recent advancements in artificial intelligence have facilitated the adoption of Machine Learning (ML) in nuclear engineering. For instance, deep convolutional neural networks have been successfully applied to predict the cycle maximum pin power peaking factor and relative fuel assembly powers in Pressurized Water Reactors with high fidelity [7]. Extending to MSRs, previous studies utilized machine learning algorithms to optimize single-channel designs, demonstrating that surrogate models can accurately predict key performance metrics such as k_{eff} and conversion ratio while drastically reducing the computational cost compared to Monte Carlo simulations [8]. Furthermore, in the field of nuclear safety analysis, machine learning algorithms have proven effective in predicting critical safety parameters, such as peak cladding temperature during loss of coolant accidents, thereby significantly reducing the computational burden of transient simulations [9]. These approaches confirm that data-driven surrogate models can effectively replace expensive physics-based simulations, accelerating the design search process.

Despite these developments, the application of advanced optimization frameworks to the MSFR utilizing liquid fuel salt remains relatively unexplored. Most existing studies on MSFRs rely on parametric sensitivity analyses or manual trial-and-error approaches [10]. Furthermore, balancing the trade-off between strict criticality safety constraints and fuel breeding performance requires a sophisticated constraint-handling mechanism. In this context, this study proposes a surrogate-based optimization procedure to determine the optimal core configuration of an MSFR. We employ a Multi-Layer Perceptron (MLP) as a fast and accurate surrogate model to approximate the complex non-linear relationships between design variables (e.g., core radius, fuel composition, temperature) and neutronic performance indicators. By coupling this surrogate with a Genetic Algorithm (GA) integrated with a penalty function, we aim to maximize the Breeding Ratio (BR) while ensuring criticality safety. This approach offers a computationally efficient workflow that bridges the gap between OpenMC simulations and practical engineering design optimization.

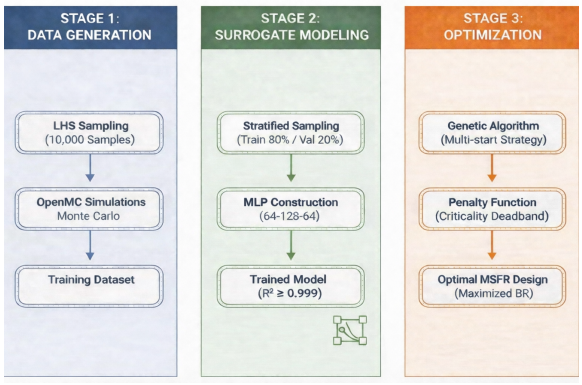

Fig. 1. Integrated framework for surrogate-based core design optimization.

2. Methods

In this study, a surrogate-based optimization procedure was employed to determine the optimal core

design of a MSFR. The optimization workflow is implemented through three sequential stages: (1) generation of a training dataset via OpenMC simulations using Latin Hypercube Sampling (LHS), (2) construction of a MLP surrogate model trained with a stratified sampling technique and (3) search for optimal design parameters using a GA integrated with a penalty function for criticality constraint-handling.

2.1 Data Generation

Neutronic simulations were performed using the OpenMC transport code to generate the training dataset. The reference core configuration adopted in this study is based on the MSFR EVOL benchmark design [11], which represents a 3,000 MWth fast-spectrum reactor. The reference reactor model is a single-channel cylindrical geometry representing a typical MSFR core fueled with the LiF-ThF_4-UF_4 system. To account for thermal feedback effects, the fuel salt density (ρ) was modeled as a function of temperature (T) using Eq. (1) [12].

Eq. (1) $\rho(T) = 4094 - 8.82 \cdot (T - 1008)$ [kg/m³]

To ensure efficient coverage of the high-dimensional design space while minimizing computational cost, a comprehensive database comprising 10,000 samples was constructed using LHS. Each design variable was stratified into 10,000 equal-probability bins, and one sample was drawn from each bin to ensure uniform marginal distributions. The selection of the five design variables and their respective search spaces, as summarized in Table I, was driven by the physical, material, and operational constraints of the MSFR. The operational temperature is bounded to maintain fuel salt liquidity. The limits for ^{233}U enrichment and the thorium blanket fraction were defined to explore the critical trade-off between ensuring core criticality and maximizing fertile-to-fissile conversion. Furthermore, the core radius was constrained to balance the neutronic advantage of reduced neutron leakage—which inherently improves the breeding ratio—against the practical spatial and economic limitations of a compact reactor footprint. Finally, the B_4C reflector thickness was set to provide adequate neutron shielding without excessively increasing the overall reactor volume.

Table I : Range of design parameters for optimization

	Reference	Range
Temp(K)	900	873 ~ 1073
^{233}U(%)	2.5	1.5 ~ 3.5
Th_blanket(%)	22.5	15 ~ 25
Radius(cm)	112.75	100~130
B_4C(cm)	20.0	10~30

2.2 MLP Model Development

A MLP was developed to serve as a fast surrogate model for the OpenMC code. The network architecture consists of an input layer (5 nodes), three hidden layers (64-128-64 neurons) with batch normalization and ReLU activation functions, and an output layer (3 nodes).

We applied a stratified sampling technique based on the distribution of k_{eff}. The dataset was divided into training (80%) and validation (20%) sets while preserving the statistical distribution of k_{eff} in both subsets. All input and output variables were standardized using a *Scikit-learn StandardScaler* to facilitate stable gradient descent.

The model training was driven by minimizing the Mean Squared Error (MSE) loss function via the Adam optimizer with a learning rate of 0.001. The training process was executed for 200 epochs. To implement the stratified sampling effectively, the range of k_{eff} was discretized into 50 bins, ensuring that the data distribution was preserved across the training and validation sets. In addition, the Mean Absolute Error (MAE) and R^2 were monitored to evaluate the surrogate model performance.

2.3 GA Optimization

The optimization problem was defined to maximize the BR while strictly satisfying the criticality constraint. The GA was employed as the optimizer. To handle the strict criticality constraint within the GA, a penalty function with a deadband was introduced. The fitness function used in the GA is defined as shown in Eq. (2). To ensure the design satisfies the strict criticality constraint, the function incorporates a multi-stage penalty mechanism that accounts for deviations from the target k_{eff} range. Specifically, a penalty is applied to solutions that deviate from the target k_{eff} lower bound of 0.983 (k_{low} = 0.983) or exceed the safety threshold of 1.05(k_{high} = 1.05), effectively guiding the algorithm to prioritize designs within the acceptable criticality margin while maximizing the BR.

$$\text{Eq. (2) Fitness} = \text{BR} - \omega \left(\frac{d}{k_{high} - k_{low}} \right)^2$$

Where ω is the penalty weight (set to 2.0), and d represents the deviation from the allowable range shown in Eq.(3)

$$\text{Eq. (3) } d = \begin{cases} k_{low} - k_{eff} & if\ k_{eff} < k_{low} \\ k_{eff} - k_{high} & if\ k_{eff} > k_{high} \\ 0 & otherwise \end{cases}$$

Finally, to overcome the stochastic nature of the GA and avoid local optima, a multi-start strategy was adopted. The optimization process was repeated for 50

independent runs. The final optimal design candidates were validated using OpenMC simulations to confirm their physical feasibility.

3. Results

3.1 Data distribution

The LHS-generated dataset for 10,000 samples showed that while k_{eff} and BR covered wide ranges (0.71–1.19 and 0.78–1.91, respectively), the fast neutron fraction remained remarkably stable with a mean of 0.314 and a narrow standard deviation of 0.012. This spectral invariance indicates that the MSFR core inherently maintains a robust fast-spectrum regime across the defined design space, rendering active spectral optimization redundant. Furthermore, fast neutron fraction was excluded from the optimization process because spectral constraints are implicitly addressed through the optimization of k_{eff} and BR, given its strong parametric coupling with ^{233}U enrichment. By streamlining the objectives in this manner, the optimization framework could focus more effectively on maximizing breeding performance while ensuring criticality safety.

3.2 Validation of Surrogate Model

The predictive performance of the MLP surrogate model was evaluated using a held-out test set comprising 20% of the total samples, constructed via stratified sampling based on k_{eff} distributions. Table II summarizes the key statistical metrics. The model demonstrated exceptional accuracy across all target variables. Specifically, the R^2 score for k_{eff} reached 0.9994. The BR prediction also exhibited high fidelity with an R^2 of 0.9992 and an MAE of 6.47×10^{-4}.

Table II : Predictive performance of the surrogate model

	k_{eff}	BR
R^2	0.9994	0.9992
MSE	9.62×10^{-6}	7.47×10^{-5}
MAE	250 pcm	6.47×10^{-4}

3.3 Optimization Trajectory

As depicted in Fig. 2, Run #15 achieved the highest fitness score among the 50 independent GA runs. Also Fig. 3 illustrates the optimization history of the best-performing run. Furthermore, the Fig. 3 demonstrates convergence to a fitness score of 1.137 after approximately 700 steps. The convergence process exhibits two distinct phases: (1) feasibility search and (2) performance maximization. In the initial phase, the algorithm explored the design space with significant fluctuations in criticality. Early solutions exhibited states violating the safety margin. However, the penalty function effectively guided the population toward the feasible region.

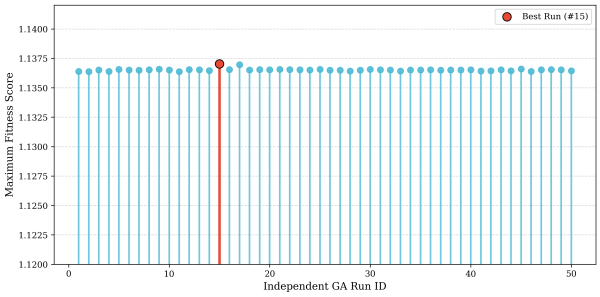

Fig. 2. Distribution of fitness scores across 50 independent GA runs.

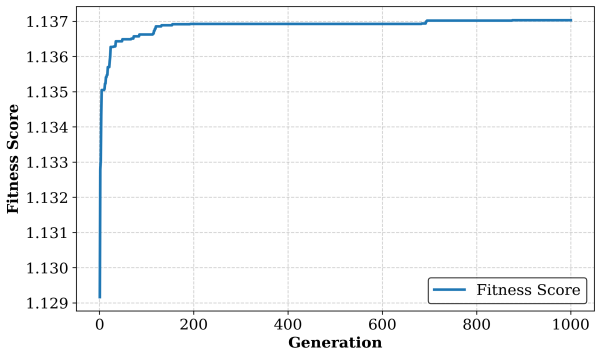

Fig. 3. Optimization trajectory for the best-performing run (Run #15).

4. Discussions

It is noteworthy that the MLP surrogate model maintained $R^2 = 0.9994$ for k_{eff} and $R^2 = 0.9992$ for BR. The neural network successfully learned the complex non-linear mapping between geometric/thermal design variables and neutronic performance indicators.

The genetic algorithm converged toward a design with a relatively large core radius (129.84 cm). From a reactor physics perspective, the preference for a larger radius can be attributed to the minimization of neutron leakage. In a fast reactor spectrum, where the neutron mean free path is long, reducing leakage demonstrates the effective way to improve neutron economy. Although this result aligns with the general expectation that an increased radius leads to a higher BR, as shown in Table 3, the other four variables were adjusted to prevent k_{eff} from deviating excessively. This indicates that the reactivity increase caused by the expanded radius was offset by changes in the other variables, thereby finding a valid solution within the given safety margins.

Table III compares the reference and optimized core parameters. The predicted k_{eff} of the optimized core was 0.98047, and the OpenMC verification resulted in 0.98185, which satisfies the criticality safety constraint within the predefined margin. The final optimized design achieved an OpenMC-validated BR of 1.1487.

Specifically, as the core radius was expanded from 112.75 cm to 129.84 cm and the B_4C reflector thickness was increased from 20.0 cm to 29.97 cm to aggressively reduce neutron leakage, the system's overall reactivity inherently increased. To offset this excess reactivity and strictly satisfy the criticality constraint, the optimizer reduced the ^{233}U enrichment from 2.50 % to 2.40 %. Concurrently, the operational temperature was lowered from 900 K to 874.76 K to increase the fuel salt density, enhancing the overall neutron economy. Interestingly, the thorium blanket fraction was reduced to its lower bound of 15.00 % (from 22.5 %); this reduction mitigates excessive parasitic absorption in the blanket region, helping to maintain core criticality even with a lower fissile inventory. Despite the lower thorium concentration, the significantly enlarged core volume and improved neutron reflection successfully maximized the overall fertile conversion, ultimately achieving a superior BR of 1.1487.

As the optimization progresses, the algorithm systematically navigates the trade-off between the criticality constraint and breeding performance. The lower bound of the target keff was intentionally set to 0.983 because the OpenMC validation of the reference reactor yielded a comparable value of approximately 0.98328. Specifically, the population evolves to minimize keff towards this lowest permissible boundary while simultaneously maximizing the BR.

Table III : Comparison of the core parameters

	Reference core	Optimized core
K_{eff}	0.98328	0.98185
BR	1.1264	1.1487
Radius(cm)	112.75	129.84
Temp(K)	900	874.76
^{233}U(%)	2.5	2.40
Th_blanket(%)	22.5	15.00
B_4C(cm)	20.0	29.97

5. Conclusions

This study successfully demonstrated the applicability of a surrogate-based optimization procedure for the design optimization of MSFR. By coupling high-fidelity OpenMC simulations with a deep learning-based surrogate model, we established a computationally efficient optimization workflow that bypasses the prohibitive costs of direct Monte Carlo iterations.

The proposed framework identified an optimal core configuration that strictly satisfies the criticality safety constraint while achieving a BR of 1.1487. Compared to the reference design, the optimized core achieved a significant improvement in the breeding ratio (from 1.1264 to 1.1487). While the OpenMC-validated k_{eff} (0.98238) showed a marginal decrease from the reference value (0.98185) due to the surrogate model's inherent prediction error, it effectively maintained the required criticality safety level. Consequently, the final design ensures both operational safety and a positive breeding gain.

Despite these results, this study is subject to certain limitations that warrant further investigation. First, the MAE of 250 pcm for k_{eff} indicates that there is room for improvement in terms of neutronic accuracy. Therefore, future work will need to enhance this precision by adopting deeper, more sophisticated MLP models or exploring other neural network architectures. Second, the current optimization was performed based on the beginning of cycle state, without accounting for the time-dependent evolution of the fuel composition. This static initial assessment explains the slightly subcritical k_{eff} observed in the optimized core, which is expected to be compensated by the high breeding ratio over a full operational cycle. Third, the thermal-hydraulic feedback was treated using a simplified temperature assumption rather than a fully coupled fluid dynamics model. Finally, the optimization framework considered a limited set of design variables and performance indicators, without simultaneously incorporating all relevant physical, material, and operational parameters that may influence reactor behavior. Therefore, the current results should be interpreted as a preliminary design optimization under constrained parameter settings rather than a fully comprehensive multi-parameter assessment.

To address these limitations and further validate the robustness of the proposed framework, future research will focus on a multi-faceted approach. We intend to conduct a comprehensive sensitivity analysis to quantitatively explain the complex correlations between geometric decision variables and neutronic stability constraints. Building upon this understanding, the optimization framework will be extended to incorporate high-fidelity multi-physics simulations. This advancement will enable the precise evaluation of performance over the entire reactor cycle and ensure thermal-hydraulic safety under realistic operating conditions.

REFERENCES

[1] IEA. (2022). Nuclear Power and Secure Energy Transitions: From Today's Challenges to Tomorrow's Clean Energy Systems. International Energy Agency.

[2] U.S. DOE Nuclear Energy Research Advisory Committee & Generation IV International Forum. (2002). A Technology Roadmap for Generation IV Nuclear Energy Systems (GIF-002-00).

[3] Serp, J., Allibert, M., Beneš, O., Delpech, S., Feynberg, O., Ghetta, V., Heuer, D., Holcomb, D., Ignatiev, V., Kloosterman, J. L., Luzzi, L., Merle-Lucotte, E., Uhlíř, J., Yoshioka, R., & Zhimin, D. (2014). The Molten Salt Reactor (MSR) in generation IV: Overview and perspectives. Progress in Nuclear Energy, 77, 308-319.

[4] Aufiero, M., Brovchenko, M., Cammi, A., Clifford, I., Geoffroy, O., Heuer, D., Hursin, M., Leppänen, J., Merle-

Lucotte, E., Munoz-Cobo, J. L., Perrot, L., Rochman, D., Wang, K., & Rubiolo, P. (2014). Calculating the effective delayed neutron fraction in the Molten Salt Fast Reactor: Analytical, deterministic and Monte Carlo approaches. Annals of Nuclear Energy, 65, 78-90.
[5] Romano, P. K., & Forget, B. (2013). The OpenMC Monte Carlo particle transport code. Annals of Nuclear Energy, 51, 274-281.
[6] Los Alamos Scientific Laboratory Group TD-6. (1978). *MCNP—A general Monte Carlo code for neutron and photon transport* (Report No. LA-7396-M). Los Alamos Scientific Laboratory.
[7] Nam, Y., & Shim, H. J. (2023). Development of deep convolutional neural network for prediction of cycle maximum pin power peaking factor in pressurized water reactor. Annals of Nuclear Energy, 194, 110083.
[8] M. Turkmen, G.J.Y. Chee, K.D. Huff, Machine learning application to single channel design of molten salt reactor, Annals of Nuclear Energy 161 (2021) 108409.
[9] Sallehhudin, W., & Diab, A. (2021). Using machine learning to predict the fuel peak cladding temperature for a large break loss of coolant accident. Frontiers in Energy Research, 9, Article 755638.
[10] Nagy, K., Kloosterman, J. L., Lathouwers, D., & van der Hagen, T. H. J. J. (2011). New breeding gain definitions and their application to the optimization of a molten salt reactor design. Annals of Nuclear Energy, *38*(2-3), 601-609.
[11] CNRS. (2014). EVOL *(*Project n°249696*)* final report. European Commission.
[12] Laureau, A., Bellè, A., Allibert, M., Heuer, D., Merle, E., & Pautz, A. (2022). Unmoderated molten salt reactors design optimisation for power stability. Annals of Nuclear Energy, *177*, 109265.