Transactions of the Korean Nuclear Society Spring Meeting

Jeju, Korea, May 7-8, 2026

Machine Learning-Based Optimization and Response Surface Analysis of Primary System Controller Gains of i-SMR

Taeho Roh, Semin Joo, Jeong Ik Lee*

Department of Nuclear and Quantum Engineering, Korea Advanced Institute of Science and Technology, 373-1 Guseong-dong Yuseong-gu, Daejeon 305-701, Republic of Korea

** Corresponding author: jeongiklee@kaist.ac.kr*

Keywords: machine learning, optimization, response surface

1. Introduction

One of the key design objectives of the i-SMR is to realize soluble boron-free (SBF) operation. By eliminating the need to adjust boric acid concentration in the coolant, SBF operation can secure a strong negative moderator temperature coefficient (MTC), thereby significantly enhancing the reactor's inherent safety. It also offers the potential to improve economic competitiveness and simplify plant systems through substantial streamlining of the Chemical and Volume Control System (CVCS). However, because fine reactivity control via boron concentration adjustment is no longer available, a large MTC can make the primary-system thermal–hydraulic conditions more sensitive to disturbances during transients such as turbine power changes. Therefore, to maintain system stability under power variations, precise NSSS control logic design and advanced control capability are essential.

The Nuclear Steam Supply System (NSSS) of a nuclear power plant consists of major components such as the reactor and steam generator, and produces high-temperature, high-pressure steam to generate turbine torque and, ultimately, electricity. In general, the NSSS is designed to accommodate small load variations and disturbances under automatic operating modes and to respond appropriately. In particular, during various transient conditions—such as load-following operations and feedwater-system disturbances—key NSSS variables (e.g., pressure, level, flow rate, and temperature) exhibit strongly coupled dynamics. Therefore, an appropriate design of the NSSS control system is essential not only to achieve satisfactory load-following performance but also to secure operating margin and to prevent unnecessary actuations of reactor safety systems.

However, tuning control gains (e.g., PI gains) to reflect realistic plant dynamics often requires repeated evaluations using system thermal-hydraulic simulation codes, which imposes substantial computational cost and long execution time, thereby limiting systematic optimization. To mitigate this burden, surrogate-model-based optimization has been increasingly adopted, where a surrogate model is trained on data generated from thermal-hydraulic system codes and then used to perform optimization efficiently. Nevertheless, to establish confidence in the obtained optimum, the surrogate-derived solution should be re-validated using the original system analysis code.

This study considers a Performance Related Design Bases Event (PRDBE) case involving a linear change in turbine power and derives optimal controller gains by applying both genetic-algorithm (GA) and gradient-descent (GD)-based optimization using a Multi-Layer Perceptron (MLP) -based surrogate model. The optimal solutions are then re-validated with MARS-KS, and the database-best point, surrogate-based optimum, and MARS-KS validation results are compared and analyzed. In addition, the surrogate model is used to investigate the characteristics of the proportional-integral (PI) gain response surface.

2. Methodology

2.1 Machine learning(surrogate) model Development

In this study, a representative PRDBE was selected to evaluate the performance of the NSSS control system: turbine power ramp change event at a rate of 5% per minute (100–20–100% ramp change). Three NSSS control subsystems were considered in this work: the Pressurizer Pressure Control System (PPCS), which regulates pressurizer pressure within the operating range using heater power through PI controllers, the Pressurizer Level Control System (PLCS), which maintains pressurizer water level within the operating range of by PI controllers of charging and letdown flow rates, and the Feedwater

Control System (FWCS), which regulates feedwater flow rate via PI controllers based on the main steam pressure setpoint.

Table 1. Range of control gain

System	Control gain	Max	Min
PPCS	P_heater	20.0	0.02
	I_heater	0.01	1.0e-5
PLCS	P_charging	50.0	0.05
	I_charging	0.01	1.0e-5
	P_letdown	50.0	0.05
	I_letdown	0.01	1.0e-5
FWCS	P_fw	50.0	0.05
	I_fw	0.01	1.0e-5

As summarized in Table 1, the range of PI gains were defined for the heater, charging, letdown, and feedwater. To efficiently sample the resulting gain space, maximin Latin hypercube sampling (maximin-LHS) was applied to generate a total of 10,000 gain combinations. For each combination, a MARS-KS transient simulation was performed to construct a database (DB). The DB consists of inputs (eight PI gain combinations) and output (the performance index: J) [1] (see Fig. 1).

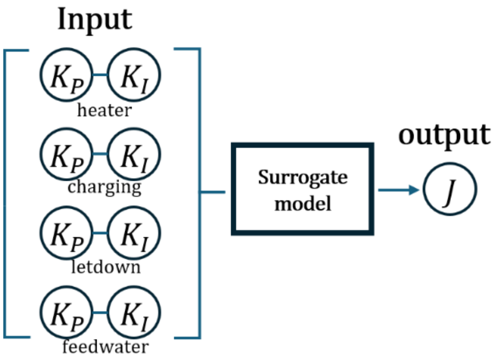

Fig. 1. Conceptual diagram of the surrogate model

For evaluating the control system, the performance index J was introduced. J accounts for not only the error between the control variables and their setpoints but also a penalty term associated with the control inputs (Eq. 1). In this formulation, the x^TQx term represents the error-related cost with respect to the setpoints, while the u^TRu term represents the control-effort penalty. The state vector x includes the errors of pressurizer pressure, pressurizer water level, and main steam pressure, and the control-input vector u includes $\dot{Q}_{heater}$, $\dot{m}_{charging}$, and $\dot{m}_{letdown}$. The weighting matrices Q and R are determined based on Bryson's rule (Eqs. 2, 3) [2], and the corresponding terms are summed to form the performance index J. A smaller value of J indicates that the controlled variables converge to their setpoints more quickly and stably.

$$J = \int_0^\infty [x^TQx + u^TRu]dt \qquad (Eq. 1)$$

$$Q_{ii} = \frac{1}{maximum\ acceptable\ value\ of\ x_i^2} \qquad (Eq. 2)$$

$$R_{ii} = \frac{1}{maximum\ acceptable\ value\ of\ u_j^2} \qquad (Eq. 3)$$

An MLP-based surrogate model was then trained on the DB to enable rapid prediction of the performance index J for a given set of PI gains. As shown in Table 2, the predictive performance of the surrogate model was assessed using metrics such as mean absolute percentage error (MAPE) and root mean squared error (RMSE), confirming that the model provides sufficient accuracy for subsequent optimization and response surface analysis.

Table 2. performance of surrogate model

MLP surrogate model	MAPE(%)	RMSE
	2.95	22.67

2.2 Optimization of control gains

Machine learning-based optimization consists of (i) rapidly evaluating the performance index J using the surrogate model, (ii) searching for the optimal combination of control gains with an optimization algorithm, and (iii) re-validating the obtained solution by applying it to MARS-KS. In this study, a Genetic Algorithm and a Gradient Descent method were adopted as the optimization algorithms.

Gradient Descent method

The gradient descent (GD) is an optimization method that iteratively updates design variables in the direction that decreases an objective function, using local gradient information. At the current point, the gradient of the objective function is evaluated, and the variables are moved in the opposite direction by an appropriate step size to search for an optimal solution.

Because GD directly exploits derivative information, it can efficiently reflect the local structure of the objective function. In particular, for convex problems, convergence to the global optimum can be guaranteed under proper step-size control. For non-convex or multi-modal response surfaces, however, GD may converge to different local optima depending on the initial point and the choice of step size. Nevertheless, GD remains an effective optimization approach due to its fast computational speed and its ability to achieve precise convergence in the vicinity of an optimum. In this study, GD was performed from multiple initial points to mitigate convergence to local optima.[3]

Genetic Algorithm

The genetic algorithm (GA) is a population-based search methodology that evolves multiple candidate solutions, in contrast to gradient descent (GD), which updates a solution using gradient information of the objective function. GA relies solely on fitness values returned by the objective function for evaluation. This black-box evaluation characteristic is particularly advantageous for nonlinear and high-dimensional response surfaces—such as the eight-dimensional PI gain space considered here—where the objective landscape would be complex and may exhibit multi-modality. By exploring the search space through probabilistic sampling rather than being constrained by local gradients, GA is generally less sensitive to initial point and is less prone to being trapped in local optima.

In practice, GA proceeds through the interaction of genetic operators governed by probabilistic transition rules over successive generations. The reproduction (selection) operator increases the survival probability of superior solutions according to their fitness, thereby driving convergence toward promising regions of the search space (exploitation). Meanwhile, the crossover operator recombines the parameter vectors of selected parent individuals to generate offspring, enabling the exploration of new regions by inheriting and mixing favorable traits (exploration). In addition, the mutation operator randomly perturbs an individual with a low probability to maintain population diversity and mitigate premature convergence. In this study, the eight PI gains are treated as genes, and GA is applied to identify a PI gain combination that minimizes the performance index. [4]

2.3 Response surface analysis

A response surface represents the functional landscape of an objective function with respect to changes in design variables, and it is useful for interpreting optimization results and assessing the suitability of an optimization strategy. In particular, in multivariable controller-gain tuning problems, the effects of individual gains on the performance index can be nonlinearly coupled. Therefore, response-surface analysis can be used to identify differences in variable sensitivity, characterize the landscape geometry, and examine the potential existence of multiple local optima. Such geometric information also provides a basis for evaluating where a given optimization algorithm converges within the search space and how robust the resulting solution is.

In this study, the response surface of the performance index J over the eight-dimensional PI gain space was computed using the trained surrogate model. Because direct visualization of the full high-dimensional design space is not feasible, a two-dimensional slice-based analysis was performed. Response surfaces were constructed for loop-wise PI gain pairs (P_{heater}, I_{heater}), $(P_{charging}, I_{charging})$, $(P_{letdown}, I_{letdown})$, and (P_{fw}, I_{fw}), while the remaining six gains were fixed at a reference point. The reference point was defined as the DB-best gain set that yields the minimum J in the DB. The resulting surfaces were visualized using 2D contours and 3D surface plots to facilitate examination of the structural characteristics of the response surface.

3. Results and Discussions

3.1 Comparison of optimal control gains

Table 3. Comparison of optimized PI gain combinations

Optimal gain	DB	GA	GD
P_heater	19.69	20.00	20.00
I_heater	9.46e-3	1.00e-2	1.00e-2
P_charging	35.57	50.00	49.98
I_charging	7.84e-3	1.00e-5	1.10e-5
P_letdown	49.48	50.00	50.00
I_letdown	3.96e-3	1.00e-5	1.10e-5
P_fw	10.62	10.76	10.76
I_fw	4.52e-3	1.00e-5	1.20e-5

Table 3 compares the optimal point within the PI gain combinations from DB generated via maximin-LHS with the PI gain combinations derived from GA and GD-based optimizations. Both GA and GD showed significant changes in gain combinations compared to the DB optimal point. In particular, they exhibited a convergence trend toward the lower bounds for the integral gains of the charging, letdown and feedwater. In contrast, proportional gains generally tended to converge toward their upper bounds for the heater, charging, and letdown. I_heater converged to approximately 0.01 for both methods, and P_fw showed a slight increase from the DB optimal point (10.62) to 10.76 for both GA and GD.

In summary, surrogate-based optimization was found to explore new optimal combinations by adjusting the integral and proportional gains of certain control loops, rather than performing simple fine-tuning around the DB optimal point. The final gains of GA and GD are overall at a similar level, suggesting that both methods explored similar high-performance regions within the same performance landscape.

3.2 Improvement of performance index and MARS-KS re-validation results

Table 4. Summary of performance index

	Performance index (J)
Optimal J (DB)	56.43
GA result	50.47
MARS validation (GA)	53.60
GD result	50.48
MARS validation (GD)	53.44

Table 4 compares the DB optimal point, surrogate optimal solutions (GA/GD), and the results of re-validating each optimal solution with the MARS-KS. [5] The performance index J for the DB optimal point is 56.43. Based on the surrogate model, J was predicted as 50.47 for GA and 50.48 for GD, indicating improved performance compared to the DB for both methods.

However, as the reliability of surrogate optimization must be confirmed through re-validation using the system analysis code, the GA and GD optimal solutions were each applied to MARS-KS to calculate the performance index. According to the MARS-KS re-validation results, the J of the GA optimal solution was evaluated as 53.60, and that of the GD optimal solution as 53.44. Both cases confirmed that performance improvement over the DB optimal point was maintained. This demonstrates that the solutions derived from surrogate-based optimization provide meaningful improvements even by actual simulation standards.

The MARS-KS re-validation performances of GA and GD are close to each other, showing that both methods achieved similar levels of improvement. Meanwhile, the MARS-KS re-validation values were somewhat higher than the surrogate-predicted values, indicating that the surrogate model predicted J is slightly optimistically near the optimal point. The prediction error rates were calculated as 5.84% for GA and 5.55% for GD, with both methods showing a discrepancy within approximately 6%. This discrepancy may be attributed to regression errors based on a limited database, the movement of the search into relatively sparse regions of the training distribution during optimization, and non-linear interactions within the system. Nevertheless, these results suggest that surrogate optimization provides solutions that maintain improvement during simulation-based re-validation, confirming that both GA and GD can be utilized as effective optimization tools.

3.3 Response surface analysis

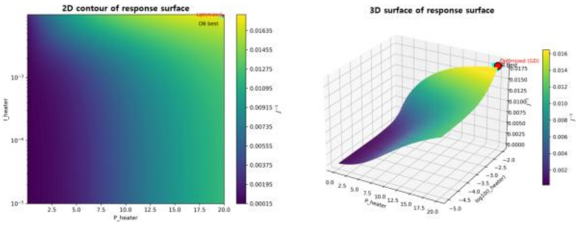

Fig. 2. 2D contour and 3D surface of response surface for (P_{heater}, I_{heater})

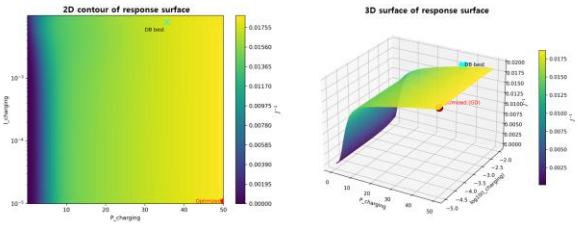

Fig. 3. 2D contour and 3D surface of response surface for $(P_{charging}, I_{charging})$

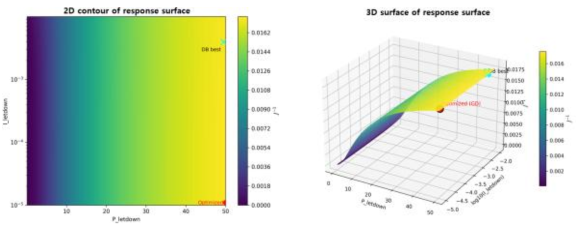

Fig. 4. 2D contour and 3D surface of response surface for $(P_{letdown}, I_{letdown})$

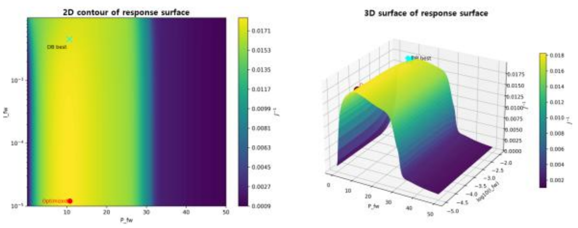

Fig. 5. 2D contour and 3D surface of response surface for (P_{fw}, I_{fw})

Figures 2–5 present the surrogate-predicted two-dimensional and three-dimensional response surfaces of the reciprocal of the performance index J, obtained by sweeping each loop-wise PI gain pair while fixing the remaining gains at the DB-best gain set. The DB-best reference point is indicated by an "X" marker, and the optimized point is indicated by an "O" marker. Overall, variations along the I gain direction are relatively mild, whereas there exist regions where J changes sharply with the P gain, which helps explain why the optimization tends to converge to specific ranges of P gains.

For the (P_{heater}, I_{heater}) surface, a steep region is formed at low P_{heater} where J increases significantly, and J decreases rapidly as P_{heater} increases. This trend is consistent with the optimization results, in which relatively large values of P_{heater} are selected. The $(P_{charging}, I_{charging})$ and $(P_{letdown}, I_{letdown})$ surfaces exhibit comparatively broad flat regions, with pronounced increases in J only near specific boundary regions. Finally, for the (P_{fw}, I_{fw}) surface, J tends to increase when P_{fw} departs from a certain range, which is consistent with the optimization behavior that avoids excessively large P_{fw} and converges to a relatively constrained region.

Across the four PI gain pairs, the response surfaces are generally smooth, suggesting a reduced likelihood of convergence to distinct local optima. This supports the observation that GA and GD, despite employing different optimizing mechanisms, can yield solutions with similar performance index. Overall, the response-surface analysis provides interpretive support for the convergence behavior and the validity of the GA- and GD-based optimization results.

However, as shown in Table 3, the result of two different optimization method do not show exact same PI gain combination. It might be because the eight PI gains do not act independently in the closed-loop NSSS dynamics; instead, a gain change in one control loop can simultaneously alter the tracking errors and control-effort demands of other loops, inducing cross-coupled interactions (a “combined effect”). As a result, even if the response surface appears smooth in a two-dimensional slice, the full eight-dimensional landscape can exhibit tilted valley/ridge directions or multiple near-optimal combinations, meaning that different gain sets may yield similar J values. Therefore, although GA and GD minimize the same objective function, differences in their search trajectories and convergence mechanisms can lead them not to converge to exact identical points within such near-optimal regions.

4. Summary and Conclusions

In this study, surrogate model-based optimization was applied to PI gain tuning of the Primary control system for a PRDBE turbine power ramp change scenario. Using an surrogate model as a fast evaluator, two optimization approaches—genetic algorithm (GA) and gradient descent (GD)-based optimization—were implemented and compared, and their candidate optima were re-validated using MARS-KS to quantify surrogate-to-simulation consistency. In addition, surrogate model-based response surface analysis was performed through two-dimensional slices of the eight-dimensional gain space, providing insight into the response surface characteristics that govern convergence behavior and the resulting gain selections. The results indicate that both GA and GD can identify gain sets that improve the performance index relative to the DB-best point, while the MARS-KS re-validation highlights the importance of surrogate fidelity when interpreting optimization outcomes. Further work will focus on further enhancing the reliability of surrogate-based optimization through active learning for data expansion near the optimal regions and the improvement of optimization algorithms considering safety constraints and uncertainties.

ACKNOWLEDGEMENT

This work was supported by the Innovative Small Modular Reactor Development Agency grant funded by the Korea Government (MSIT: Ministry of Science and ICT) (No. RS-2024-00405419)

REFERENCES

[1] S. M. Joo, and J. I. LEE, “Machine Learning-Based Prediction of Pressurizer Control Performance in i-SMR under Load-Following Conditions”, in Proc. Korean Nuclear Society Autumn Conf., Changwon, Korea, 2025

[2] A. E. Bryson, Jr., and Y.-C. Ho, Applied Optimal Control: Optimization, Estimation, and Control, Hemisphere Publishing Corporation, 1975.

[3] S. Boyd, and L. Vandenberghe, Convex Optimization, Cambridge University Press, 2004

[4] D. E. Goldberg, Genetic Algorithms in Search, Optimization, and Machine Learning, Addison-Wesley Publishing Company, 1989

[5] B. D. Jeong, et al., MARS Code Manual Volume 1: Code Structure, System Models, and Solution Methods, KAERI/TR 2812/2004, Korea Atomic Energy Research Institute, 2004.