Proc. of the European Safety and Reliability Conference (ESREL 2026)
Edited by José C. Matos, Paulo B. Lourenço, Michael Beer, Daniel Oliveira,
Edoardo Patelli, Hélder S. Sousa, Rui A. Silva and Monica Santamaria

doi: 10.3850/ESREL2026061419_esrel26-p26454-cd

Adaptive Sampling Framework for Data-Efficient AI Prediction of Parameter Trends in Small Modular Reactor Simulations

Donghee Jung

Department of Nuclear Engineering, Ulsan National Institute of Science and Technology (UNIST), Republic of Korea. E-mail: joy00@unist.ac.kr

Hyun Gook Kang

Department of Mechanical, Aerospace, and Nuclear Engineering, Rensselaer Polytechnic Institute (RPI), United States

Seung Jun Lee

Department of Nuclear Engineering, Ulsan National Institute of Science and Technology (UNIST), Republic of Korea. E-mail: sjlee420@unist.ac.kr

Small modular reactors (SMRs) require reliable monitoring and operator-support capabilities during abnormal conditions, but AI forecasting models for SMR transients require simulation data spanning malfunction severity and mitigation timing. This study proposes a hierarchical adaptive sampling framework for data-efficient generation of forecasting datasets in the IAEA integral pressurised water reactor (iPWR) simulator. Level 1 adaptively selects leakage severities from a no-action pool to diversify transient regimes and reduce redundant sampling in saturated early-trip regions. Level 2 refines isolation-valve closure timing for selected severities to densify samples near trip-margin-relevant timing bands. An LSTM model is iteratively trained on accumulated scenarios and used to score informative candidates by forecasting error. Results for a pressurizer relief-valve leakage case show that the method concentrates samples near severity and timing transition boundaries, improving regime coverage under a limited simulation budget. The framework is intended as an offline simulation-campaign strategy supporting future AI-based predictive monitoring and operator-support applications.
Keywords: Adaptive sampling, small modular reactor, iPWR simulator, parameter trend prediction, operator support, LSTM

1. Introduction

Small modular reactors (SMRs) are being developed for modular deployment, shorter construction schedules, and enhanced passive safety compared with conventional large nuclear power plants [1]. Their compact and highly integrated architectures increase the need for advanced monitoring, forecasting, and operator-support capabilities during abnormal operating conditions [1-3]. Because AI deployment in the nuclear industry remains at an early stage, verification, validation, transparency, and operational integration must be considered carefully [2].

Related work and research gap

Prior AI studies for nuclear power plant (NPP) operation have addressed event diagnosis, signal validation, transient identification, and parameter prediction. Examples include dynamic neural-network-based accident diagnosis, intelligent emergency-response support, explainable diagnostic assistants, deep-learning-based transient identification, and real-time prediction of parameter trends after operator actions [3-9]. These studies demonstrate the value of AI-assisted monitoring and forecasting, but they generally assume that an appropriate training corpus has already been prepared.

Active learning and adaptive sampling provide a principled way to reduce data-generation cost by acquiring informative samples [10,11]. In nuclear engineering, related approaches have been applied mainly to limit-surface search and probabilistic safety-margin evaluation, where simulations are concentrated near success/failure boundaries [12-14]. Our prior work also showed that error-guided adaptive sampling can reduce redundancy in LOCA scenarios using the compact nuclear simulator [15].

A gap therefore remains between AI-based NPP forecasting studies, which often rely on fixed datasets, and nuclear adaptive-sampling studies, which often focus on limit-surface identification. Comparatively limited attention has been paid to how a constrained simulation budget should be allocated across both malfunction severity and mitigation timing for multistep forecasting of SMR transient parameter trends near trip-relevant regions. The objective of this study is not to design an autonomous control algorithm, but to design an offline simulation-campaign strategy for constructing informative training data for predictive monitoring and operator support.

The main contributions are threefold: (1) a hierarchical adaptive sampling framework that separates severity selection from mitigation-timing refinement, (2) evaluation of sampling behavior using trip-margin-relevant regime coverage rather than only global MSE over a dense grid, and (3) demonstration in the IAEA iPWR simulator using a pressurizer relief-valve leakage scenario and an LSTM-based multistep forecasting model.

2. Problem Setup and Data

2.1. Simulator and Scenario Definition

Data were generated with the IAEA integral pressurised water reactor (iPWR) simulator, a basic-principles educational simulator intended to reproduce realistic plant trends for training and understanding rather than plant-specific safety analysis [16,17]. All simulations used IC#4, corresponding to 100% BOL forced circulation, with nominal conditions of approximately 150 MWt thermal power, 45 MWe generator load, 15.5 MPa primary pressure, 43% pressurizer level, 2.7 MPa steam-header pressure, and 78 kg/s feedwater flow [16,17].

The abnormal event is a pressurizer power-operated relief-valve (PORV) leakage scenario. The leakage-rate setting, denoted by s, represents malfunction severity. The mitigation action is closure of the corresponding isolation valve, which is physically consistent with the simulator design in which each relief valve is placed in series with a motorized valve to prevent a LOCA if the relief valve fails open [16].

Three target variables were selected to represent pressure, power, and inventory-related behavior: RCSPT02_TR (upper-plenum pressure), RCS_TP_MW (reactor thermal power), and RCSLT02_TR (pressurizer level). The main reactor-trip setpoints used to interpret the results include low upper-plenum pressure ($P<11.0$ MPa), low pressurizer level ($L<5.0\%$), low downcomer flow ($W<347$ kg/s), high core outlet temperature ($T>340$ degrees C), high neutron flux (>120%), high log rate (SUR>2 dpm), and high upper-plenum pressure ($P>16.4$ MPa) [16,17]. ADS actuation occurs for low-low upper-plenum pressure ($P<9.0$ MPa), low-low vessel level ($L<90\%$), high-high upper-plenum pressure ($P>17.2$ MPa), or high-high containment pressure ($P>0.019$ MPa) [17].

2.2. *Severity and timing pools*

The scenario space is two-dimensional: leakage severity and mitigation timing. For Level 1, no mitigation action is taken after malfunction onset, and the objective is to build a compact severity set spanning non-trip, transition, pre-saturation, and saturated early-trip regimes. The severity axis is discretized over s ranging from 1 to 50 using interleaved pools: s = 1.0, 1.5, ..., 50.0 for training candidates and s = 1.25, 1.75, ..., 49.75 for evaluation.

For Level 2, representative severities sev10 and sev30 are selected for timing refinement. The leakage malfunction is injected at 30 s, and the isolation-valve closure time varies from 31 to 630 s. Odd action times (31, 33, ..., 629 s) are used as training candidates, whereas even action times (32, 34, ..., 630 s) are reserved for evaluation, preserving strict scenario-level separation.

2.3. *Forecasting task and sample construction*

For mitigation-timing experiments, prediction is defined relative to the immediate post-action state. For action time a, the reference time is tref=a+1 s. The input consists of two plant snapshots separated by 30 s, [X(tref-30 s), X(tref)], and the output is a 10-minute multistep trajectory of the three-target vector [RCSPT02_TR, RCS_TP_MW, RCSLT02_TR] at 30-s intervals from tref+30 s to tref+600 s.

Each scenario is stored as a 22-min CSV trajectory sampled at 1 Hz. The time column is excluded, and all remaining simulator variables are used as input features. Sliding windows are generated only within training scenarios after the mitigation action, with a stride of 1 s and a full 10-min prediction horizon. Evaluation is scenario-level: only the single reference point at tref is evaluated for each held-out scenario.

Level 1 operates on no-action severity cases and is therefore analyzed mainly through the trip-time regimes associated with selected severities; Level 2 performs post-action timing refinement for the forecasting task.

2.4. *Preprocessing and forecasting model*

All splits are scenario-based. Inputs and targets are standardized using StandardScaler statistics computed only from training data and then applied consistently to evaluation scenarios. All experiments use a single LSTM with 128 hidden units that outputs a multistep forecast trajectory; training minimizes MSE for 300 epochs without early stopping. Architecture, preprocessing, optimization, and evaluation settings are fixed across adaptive and baseline comparisons to isolate the effect of data selection.

3. Two-level Adaptive Sampling Framework

3.1. *Overview and Design Rationale*

The proposed framework is hierarchical and self-reflective. A forecasting model is trained on the currently available scenarios; candidate scenarios are evaluated using the same scenario-level difficulty signal; the highest-scoring candidate is added to the training set; and new candidates are generated around the selected point for the next iteration. Thus, simulation effort is allocated toward informative regions of the (s,a) space rather than uniformly across the full domain.

The hierarchy reflects the physical structure of the problem. Leakage severity determines the transient regime, including whether the plant remains non-trip, enters a transition band, or quickly saturates into an early-trip regime. Mitigation timing then refines the response within a chosen severity by determining whether the action is taken far from, near, or beyond a trip-relevant boundary. Accordingly, Level 1 selects severities from a no-action pool, and Level 2 refines timing for selected severities.

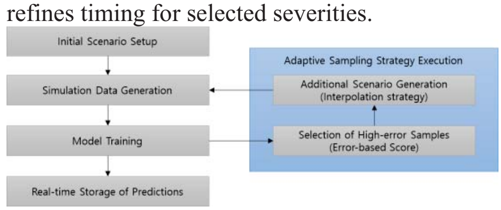

Fig. 1. Workflow of the proposed adaptive sampling framework for AI prediction in nuclear power plant simulations.

3.2. *Level 1: Severity selection from the no-action pool*

Level 1 adaptively selects severities to diversify transient regimes under a limited simulation budget. Candidate severities are evaluated using the model weakness signal, the highest-scoring severity is added to the training set, and new candidates are inserted between neighboring training points. This midpoint-based refinement follows the one-dimensional adaptive expansion strategy used in our previous work [15], here applied to the severity axis in the iPWR environment.

3.3. *Level 2: Mitigation-timing refinement*

For a fixed severity, Level 2 refines the isolation-valve closure time. Candidate timings are evaluated, the most informative timing is selected, and new timing candidates are inserted between neighboring training timings. The goal is to increase resolution near critical timing bands where small timing differences produce large changes in future pressure, power, inventory response, and trip margin.

4. Experimental Design and Evaluation

4.1. *Experimental Overview*

The experiments validate whether the two-level framework selects scenarios aligned with trip-margin-relevant forecasting. Level 1 selects severities from a no-action pool to diversify trip-time regimes. Level 2 refines mitigation timing for two representative severities in the leakage scenario to densify timing points near trip-adjacent transition regions.

4.2. *Adaptive-sampling settings and budgets*

Adaptive sampling proceeds iteratively under a fixed scenario budget. At each iteration, the current model evaluates candidate scenarios, the highest-scoring candidate is appended to the training set, and new candidates are generated around the selected point.

For Level 1, the final selected severity set reported in this paper corresponds to 10 adaptive additions, yielding 13 selected severities in total. For Level 2, timing-refinement results are shown over 10 adaptive iterations for the two representative severities, sev10 and sev30.

4.3. *Evaluation criteria focused on trip-margin-relevant coverage*

MSE is used as the training loss and internal candidate-selection difficulty signal, but the primary objective is not to minimize a single global MSE over a densely covered grid. Such a metric mainly measures interpolation across the entire domain and can favor uniform coverage. Instead, this study evaluates whether selected scenarios align with trip-margin-relevant objectives.
At Level 1, selected severities are mapped to trip-time regimes to determine whether adaptive sampling reduces redundancy in saturated early-trip regions while reallocating cases to pre-saturation regimes. At Level 2, the main criterion is whether selected action times become dense near boundary-adjacent bands where small timing differences cause large changes within the 10-minute forecasting horizon.

4.4. *Baseline construction*

For both levels, a uniform grid baseline is constructed under the same scenario budget as the adaptive method. The baseline selects severities or action times evenly across the corresponding range. All preprocessing, model architecture, training settings, and evaluation rules are kept identical to those used for adaptive sampling. The baseline therefore provides a direct comparison of how the adaptive method reallocates limited simulation budgets away from redundant regions and toward transition or boundary-adjacent regions.

5. Results

5.1. *Level 1: Severity selection behavior in the no-action pool*

The Level-1 experiment tests whether adaptive severity selection allocates limited simulation budgets toward dynamically distinct regimes rather than redundant saturated early-trip cases. The trip-time curve in Fig. 2 shows four regimes: non-trip behavior at low severity, a transition region where trip begins, a pre-saturation band where trip time changes rapidly, and a saturated early-trip band at higher severity.

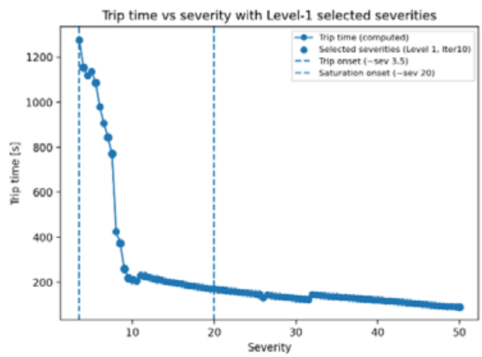

Fig. 2. Trip time versus severity derived from the no-action pool, with Level-1 selected severities overlaid to illustrate regime-focused allocation.

After initial exploration, the selected severities concentrate in the pre-saturation range rather than repeatedly sampling high-severity cases that produce similar early-trip outcomes. This behavior is consistent with the objective of increasing the diversity of trip-time regimes represented in the training set.
Figure 3 directly compares trip-time coverage between the adaptive selection and a uniform severity grid under the same budget of 13 selected severities. The uniform grid places many samples in the saturated early-trip band, whereas the adaptive method allocates 10 cases to the pre-saturation regime and only 2 to the saturated regime; the uniform grid allocates 4 and 8 cases, respectively. At iteration 10, the final adaptive severity set consists of 1.0, 4.0, 5.5, 7.0, 7.5, 8.5, 9.0, 9.5, 10.0, 11.5, 13.0, 25.5, and 50.0, confirming that most selected cases are concentrated in the pre-saturation regime rather than the saturated high-severity region.

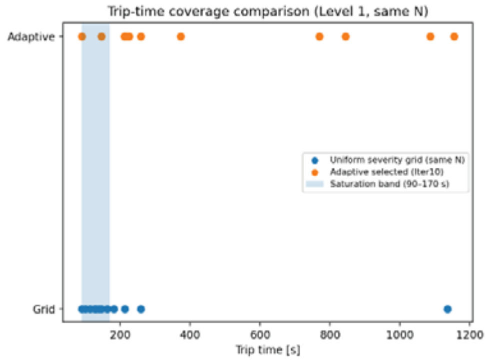

Fig. 3. Trip-time coverage comparison between adaptive Level-1 selection and a uniform severity grid baseline under the same number of selected severities; the shaded region indicates the saturated early-trip band.

These results show that Level 1 reduces redundant sampling in saturated regimes while improving coverage of dynamically diverse, trip-relevant regimes.

5.2. *Level 2: Timing refinement for the leakage mitigation action*

Level 2 refines isolation-valve closure timing to densify training cases near regions most relevant to trip margin and rapid trend changes within the 10-minute prediction horizon. Results are shown for sev10 and sev30, two representative severities that exhibit early trips.

For sev10, the selected action timings shown in Fig. 4 concentrate near a critical timing band adjacent to trip-relevant transitions rather than spreading uniformly over the full timing range.

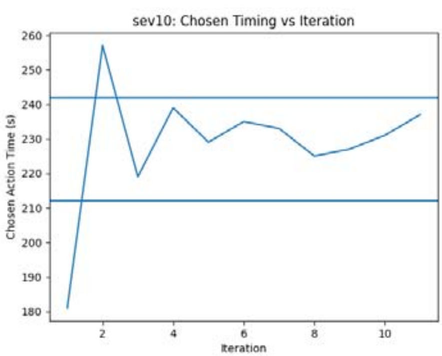

Fig. 4. Level-2 selected mitigation action timing for sev10 in the leakage scenario, showing refinement toward a critical timing region.

For sev30, Fig. 5 shows a similar refinement behavior, with selected timings becoming dense around the timing band where small action-time differences induce large changes in future trajectories and trip margin. Considering both the initially available timings and the adaptively added timings, the final training set becomes locally denser around these selected critical timing bands, indicating that Level 2 reallocates timing samples toward trip-margin-sensitive regions.

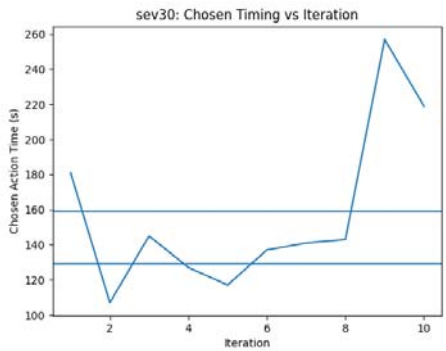

Fig. 5. Level-2 selected mitigation action timing for sev30 in the leakage scenario, showing refinement toward a critical timing region.

5.3. *Summary of Two-level Outcomes*

Taken together, Level 1 reallocates severity samples to diversify trip-time regimes and reduce saturated-region redundancy (Figs. 2 and 3), while Level 2 densifies mitigation timing near critical trip-margin-relevant bands (Figs. 4 and 5). The two-level framework therefore produces scenario-selection behavior consistent with trip-margin-aware data generation for multistep SMR parameter-trend forecasting.

6. Discussion

The results support the central premise of this paper: when the objective is informative forecasting under abnormal operation, scenario selection should emphasize transition boundaries and trip-margin-relevant regimes rather than uniformly covering the entire scenario space. In this sense, the proposed method functions as a regime-aware data-allocation mechanism, consistent with active learning and adaptive sampling [10,11].

A practical implication is that global mean MSE on a dense grid is not always aligned with the operational objective of trip-margin-aware forecasting data generation. Dense evaluation primarily measures interpolation and may favor uniform coverage. By contrast, the goal here is to capture boundary-adjacent regions where dynamics change rapidly and forecasting errors tend to concentrate, so the main claim is targeted scenario-selection behavior rather than universal dominance in global MSE.

The Level-1 results illustrate the value of avoiding redundant high-severity early-trip cases and reallocating budget to pre-saturation regimes where trip times vary widely. The Level-2 results further indicate that mitigation timing should be sampled non-uniformly because timing bands near trip-relevant transitions are both operationally important and difficult to model under a 10-minute forecasting horizon.

Several limitations remain. The experiments focus on one malfunction type and one representative mitigation action, and the framework is demonstrated in an educational simulator that is not intended for plant-specific safety analysis [16,17]. The current implementation also uses a fixed LSTM architecture and MSE-based scoring signal to isolate sampling effects. Future work should consider uncertainty-aware criteria [11,18], joint severity-timing selection, additional adaptive-sampling strategies, and broader malfunction and mitigation sets [12-14].

7. Conclusion

This paper presented a hierarchical adaptive sampling framework for data-efficient AI prediction of multivariate parameter trends in SMR simulations using the IAEA iPWR environment. The framework separates severity selection from mitigation-timing refinement and allocates limited simulation budgets toward transition boundaries and trip-margin-relevant regimes rather than uniformly sampling the full scenario space.

At Level 1, adaptive severity selection reduced redundancy in the high-severity saturated early-trip regime and increased coverage of pre-saturation severities where trip time varies widely. At Level 2, adaptive timing refinement densified training scenarios near critical closure-time bands that govern rapid trend changes and reduced trip margin. These results support adaptive sampling as a practical offline methodology for efficient forecasting-oriented dataset generation and future AI-assisted operator support in SMR operation.

Acknowledgement

This research was supported by the National Research Council of Science & Technology(NST) grant by the Korea government (MSIT) (No. GTL24031-000)

This work was supported by the Korea Institute of Energy Technology Evaluation and Planning(KETEP) and the Ministry of Trade, Industry & Energy(MOTIE) of the Republic of Korea (No. RS-2024-00401407).

References

[1] INTERNATIONAL ATOMIC ENERGY AGENCY, *Small Modular Reactors: Advances in SMR Developments 2024*, IAEA, Vienna, 2024, doi: 10.61092/iaea.3o4h-svum.

[2] INTERNATIONAL ATOMIC ENERGY AGENCY, *Considerations for Deploying Artificial Intelligence Applications in the Nuclear Power Industry*, IAEA, Vienna, 2025, doi: 10.61092/iaea.s6uy-wjt8.

[3] J. Bae, S.J. Lee, Current Progress in the Application of Artificial Intelligence for Nuclear Power Plant Operation, *Korean Journal of Chemical Engineering* 41 (2024) 2851-2870, doi: 10.1007/s11814-024-00246-7.

[4] S.J. Lee, P.H. Seong, A dynamic neural network based accident diagnosis advisory system for nuclear power plants, *Progress in Nuclear Energy* 46 (2005) 268-281, doi: 10.1016/j.pnucene.2005.03.009.

[5] J.S. Kang, S.J. Lee, Concept of an intelligent operator support system for initial emergency responses in nuclear power plants, *Nuclear Engineering and Technology* 54 (2022) 2453-2466, doi: 10.1016/j.net.2022.02.010.

[6] J.H. Park, H.S. Jo, S.H. Lee, S.W. Oh, M.G. Na, A reliable intelligent diagnostic assistant for nuclear power plants using explainable artificial intelligence of GRU-AE, LightGBM and SHAP, *Nuclear Engineering and Technology* 54 (2022) 1271-1287, doi: 10.1016/j.net.2021.10.024.

[7] I. Ramezani, N. Vosoughi, M.B. Ghofrani, Application of deep learning techniques for nuclear power plant transient identification, *Annals of Nuclear Energy* 194 (2023) 110113, doi: 10.1016/j.anucene.2023.110113.

[8] J. Bae, G. Kim, S.J. Lee, Real-time prediction of nuclear power plant parameter trends following operator actions, *Expert Systems with Applications* 186 (2021) 115848, doi: 10.1016/j.eswa.2021.115848.

[9] J.M. Kim, J. Bae, S.J. Lee, Strategy to coordinate actions through a plant parameter prediction model during startup operation of a nuclear power plant,

Nuclear Engineering and Technology 55 (2023) 839-849, doi: 10.1016/j.net.2022.11.012.
[10] B. Settles, *Active Learning Literature Survey*, Technical Report 1648, University of Wisconsin-Madison, Madison, WI, 2010.
[11] Y. Gal, Z. Ghahramani, Dropout as a Bayesian Approximation: Representing Model Uncertainty in Deep Learning, in: *Proceedings of the 33rd International Conference on Machine Learning (ICML)*, PMLR 48, 2016, pp. 1050-1059.
[12] J. Bae, J.W. Park, S.J. Lee, Limit surface/states searching algorithm with a deep neural network and Monte Carlo dropout for nuclear power plant safety assessment, *Applied Soft Computing* 124 (2022) 109007, doi: 10.1016/j.asoc.2022.109007.
[13] H. Chen, H. Wang, L. Wang, Q. Zhao, Evaluation of probabilistic safety margin of nuclear power plant based on optimized adaptive sampling method, *Annals of Nuclear Energy* 201 (2024) 110409, doi: 10.1016/j.anucene.2024.110409.
[14] H. Chen, L. Wang, H. Wang, Q. Zhao, Active learning enhanced adaptive sampling for rapid limit surface search in nuclear power plant accident scenarios, *Nuclear Engineering and Design* 446 (2026) 114589, doi: 10.1016/j.nucengdes.2025.114589.
[15] D. Jung, J. Bae, S.J. Lee, Adaptive Sampling for Data-efficient AI Prediction in NPP Simulations, *Transactions of the Korean Nuclear Society Autumn Meeting*, Changwon, Korea, 30-31 October 2025.
[16] INTERNATIONAL ATOMIC ENERGY AGENCY, *Integral Pressurised Water Reactor Simulator Manual*, Training Course Series No. 65, IAEA, Vienna, 2017.
[17] INTERNATIONAL ATOMIC ENERGY AGENCY, *Integral Pressurised Water Reactor Simulator Manual: Exercise Handbook*, Training Course Series No. 65, IAEA, Vienna, 2017.
[18] T.C.H. Nguyen, A. Diab, Uncertainty-aware prediction of Peak Cladding Temperature during extended station blackout using Transformer-based machine learning, *Nuclear Engineering and Design* 437 (2025) 113984, doi: 10.1016/j.nucengdes.2025.113984.