Transactions of the Korean Nuclear Society Spring Meeting
Jeju, Korea, May 7-8, 2026

Probabilistic Prediction of Nuclear Power Plant Transients via TimeGrad-inspired Diffusion Modeling with Operator Action Timing

Daehyung Lee*, Yongju Cho, Chungwon Seo
KEPCO E&C, Nuclear Technology Research Dept., 269 Hyeoksin-ro, Gimcheon-si, 39660, Korea
**Corresponding author: bigbro@kepco-enc.com*

Keywords : Time-series Prediction, Diffusion Model, Probabilistic Forecasting, Operator Action, Severe Accident

1. Introduction

Forecasting plant states during nuclear power plant accidents is important for operator decision support and safety. In our previous study [1], prediction accuracy improved when future operator interventions were represented by action-timing features rather than only binary on/off signals. That result showed that accident progression depends not only on the current plant state but also on when key operator actions are scheduled.

The previous Transformer-based model, however, focused on deterministic point prediction. In severe-accident scenarios, abrupt state changes and delayed interventions can produce multiple plausible future paths, so a point estimate alone is not sufficient for practical support. A probabilistic model that provides both a central trajectory and an uncertainty range is more useful in this setting.

This study extends the previous framework to a TimeGrad-inspired conditional diffusion model [2]. The same severe-accident dataset, train/validation/test split, and operator-action-aware feature design were retained. The goal was to predict containment pressure, containment hydrogen concentration, and reactor vessel water level for 120 min beyond a 20 min observation window while also estimating predictive uncertainty.

2. Methodology

2.1 Data Configuration

The dataset contains 3,000 Large Break Loss-of-Coolant Accident (LBLOCA) severe-accident simulation cases generated using the MAAP5 code [3]. The cases vary in break location and size and in the actuation timing of Safety Injection (SI), Containment Spray (CS), and Cavity Flooding (CF). The data were divided into training, validation, and test sets of 2,000, 500, and 500 cases, respectively.

Following the previous study, the first 20 min after accident initiation were used as input, and the following 120 min were predicted. The input vector contains 21 features: major process variables, scenario descriptors, operator action status, and operator action timing. The timing feature is the relative time between the scheduled action and the current time step. A future-known time covariate based on sinusoidal time encoding was also used for the diffusion model.

Table I: Experimental data configuration

Item	Value
Scenario	3,000 LBLOCA severe accident
Dataset split	Train/Validation/Test = 2,000 / 500 / 500 cases
Input window	20 minutes
Prediction horizon	Next 120 minutes
Input features	21 features: process variables, scenario descriptors, action status, and action timing
Prediction targets	Containment pressure, H2 concentration, RV water level

2.2 TimeGrad-inspired diffusion model

The proposed forecaster is a horizon-joint conditional diffusion model. The past input window is summarized by mean pooling across time and projected into a hidden context vector. Future time covariates are projected to the same hidden space and added to the context to form a conditional sequence for each forecast step. The denoising network receives the noisy target, the conditional sequence, and the diffusion-step embedding, and it is trained to predict injected Gaussian noise.

At inference, the model starts from Gaussian noise and reconstructs future trajectories through iterative reverse diffusion. Repeated sampling provides a predictive mean and quantile bands. The main hyperparameters were hidden dimension 128, diffusion steps 50, time-embedding dimension 64, batch size 128, and AdamW optimization with a learning rate of 1e-3.

2.3 Evaluation

The best checkpoint was selected by validation loss. During test-time evaluation, multiple reverse-diffusion samples were generated for each input sequence. Their sample mean was used as the point prediction, and the P10-P90 range was used as the predictive interval. Point accuracy was measured by mean absolute error (MAE) and root mean squared error (RMSE). Interval quality was measured by empirical P10-P90 coverage.

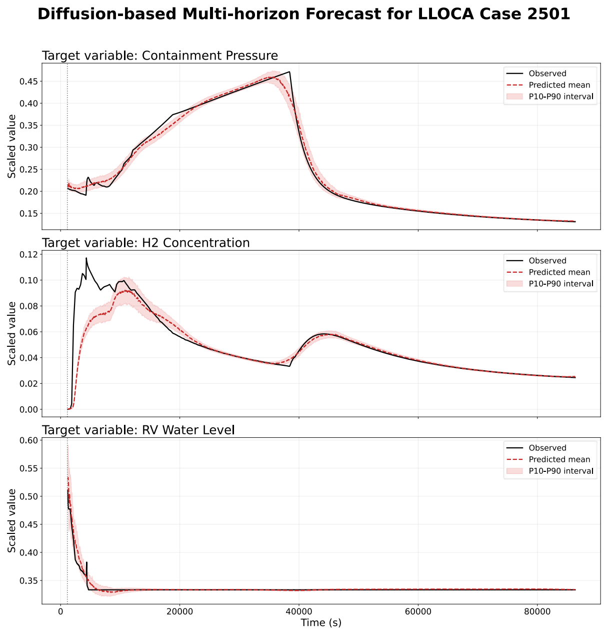

Fig. 1. Representative test case (Case 2501). Solid: actual trajectory. Red dashed: predictive mean. Shaded band: P10-P90 interval

3. Results and Discussion

The final model achieved low point-prediction error while also producing probabilistic forecasts. The test MAE and RMSE indicate that the model learned the dominant accident progression pattern well. The P10-P90 coverage, however, was lower than the nominal target of 0.80, which means that the interval width was narrower than needed for well-calibrated uncertainty. In other words, the predictive mean was accurate, but the uncertainty band still underestimated the full spread of possible outcomes.

A representative test case is shown in Fig. 1. For containment pressure, the model follows the gradual rise and the sharp drop near the major transition around 40,000 s. For H2 concentration, it reproduces the initial increase, the later decay, and the secondary hump with good agreement. For RV water level, it captures the early rapid decline and the later near-steady region. The predictive band widens around rapid transitions and becomes narrower in smoother regions, which is physically reasonable because abrupt changes are harder to predict.

The current model should be viewed as an initial diffusion benchmark rather than a fully optimized system. Better calibration may be obtained through a stronger conditional backbone, a more expressive denoiser, a refined diffusion schedule, more forecast samples, and a direct comparison with the previous Transformer baseline under the same three-target setting.

Table II: Overall performance of the diffusion model

Metric	Value
Best validation loss	0.0561
Test MAE	0.0070
Test RMSE	0.0263
P10-P90 coverage	0.6498

4. Conclusion

This study presented a TimeGrad-inspired conditional diffusion model for probabilistic forecasting of nuclear power plant transients. Using the same 3,000-case LBLOCA dataset and operator-action-aware input design as the previous study, the proposed model predicted three key variables for 120 min after a 20 min observation window. The model achieved low point-prediction error and generated useful predictive intervals through repeated reverse-diffusion sampling. Although interval calibration remains incomplete, the results show that diffusion-based probabilistic forecasting is a practical extension of the previous deterministic framework for accident-management support.

Acknowledgment

This work was supported by the Korea Institute of Energy Technology Evaluation and Planning(KETEP) and the Ministry of Climate, Energy & Environment(MCEE) of the Republic of Korea (No. 20224B10100130)

REFERENCES

[1] D. Lee, Y. Cho, C. Seo, and B. J. Kim, "Advanced Prediction of Nuclear Power Plant Transients via Operator Action-Timing Feature Engineering," Transactions of the Korean Nuclear Society Autumn Meeting, 2025.
[2] K. Rasul, C. Seward, I. Schuster, and R. Vollgraf, "Autoregressive Denoising Diffusion Models for Multivariate Probabilistic Time Series Forecasting," ICML, 2021.
[3] Electric Power Research Institute, Inc., "MAAP 5 User's Manual," 2008.