Parameter Trend Prediction Using Latent Space Forecasting for Operator Support in NPP Accidents

Donghee Jung, Junyong Bae, Seung Jun Lee

UNIST, Ulsan National Institute of Science and Technology, 50 UNIST-gil, Ulsan 44919, Republic of Korea
Joy00@unist.ac.kr

EXTENDED ABSTRACT

In emergency accident scenarios at nuclear power plants, operators must quickly assess the evolving plant status and determine appropriate mitigation actions. Real-time forecasting of key parameter trends can significantly enhance situational awareness and decision-making. While machine learning-based surrogate models have been developed for offline simulation analysis, their limited ability to process real-time signals restricts their use for operator support. To overcome this, we propose a latent space-based trend prediction framework using an LSTM Autoencoder and a decoder-assisted forecasting model. By optimizing predictions in the reconstructed signal space, the proposed approach improves the accuracy and interpretability of future trend forecasts during critical events.

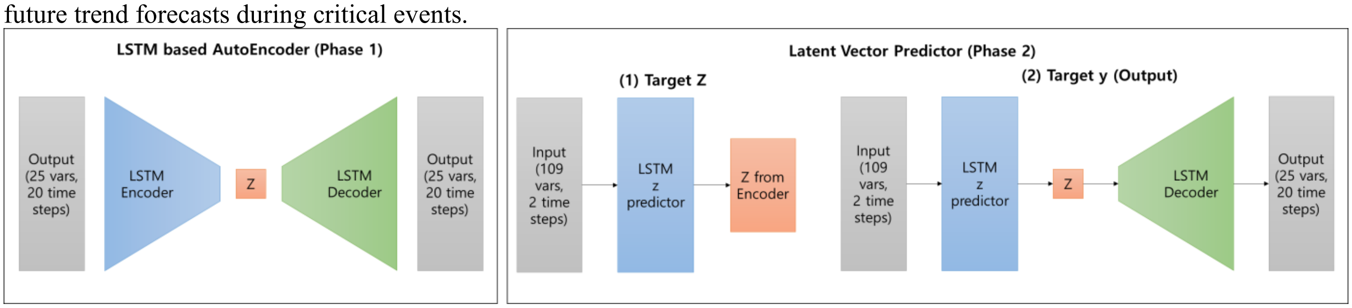

Figure 1 Overview of the proposed latent-space-based trend prediction framework. (1) latent vector prediction minimizing z error, and (2) minimizing y error.

The proposed prediction framework consists of two core components trained in two separate phases. An overview of this two-phase architecture is illustrated in **Figure 1**. In Phase 1, an LSTM-based Autoencoder (AE) compresses each 25-dimensional plant parameter sequence (20 time steps) into an 8-dimensional latent vector. The encoder is composed of seven stacked LSTM layers with decreasing units (512 → 8), each followed by batch normalization to improve convergence. The decoder mirrors this structure in reverse, expanding the latent representation back into full signal sequences. The AE is trained to minimize mean squared error between the input and reconstructed output, effectively learning a compact representation of plant dynamics. This latent space can be visualized using techniques like UMAP or t-SNE to assess the coverage of the training data and the generalizability of the model.

In Phase 2, a separate LSTM-based prediction model is trained to forecast the next latent vector using the recent plant state. This predictor receives a 2-step sequence of 743 input variables and consists of five LSTM layers (128 → 8) with dropout regularization. The predicted latent vector is then passed through the frozen decoder from Phase 1 to reconstruct the future signal trend. The entire model is optimized end-to-end using the true future sequence as ground truth. Compared to directly minimizing latent vector error (z-loss), this decoder-assisted training strategy (y-loss) encourages the predictor to capture more informative and interpretable trends in the plant behavior, leading to higher fidelity forecasting.

The framework is validated using data from the Compact Nuclear Simulator (CNS), covering scenarios such as LOCA (Loss of Coolant Accident), SGTR (Steam Generator Tube Rupture), and spurious trips. Each includes various operator mitigation actions. To determine the optimal configuration of the Autoencoder, we performed a sensitivity analysis by varying the number of LSTM layers, units, and latent dimensions. Among the tested configurations, the best result was achieved using a 14-layer structure with 512 units per layer and a latent space of 8 dimensions. This model achieved the lowest reconstruction error with a mean squared error (MSE) of 0.000235 and was used as the base for all subsequent experiments.

To evaluate model performance, both the latent-vector-based model and the full-trend model were trained on the same input data and evaluated using mean absolute error (MAE), MSE, root mean squared error (RMSE), and dynamic time warping (DTW). The results are summarized in **Table 1**. The full-trend prediction model consistently outperformed the latent-vector-based model across all evaluation metrics.

Table 1. Prediction Errors by Model Type

Model Type	MAE	MSE	RMSE	DTW
Latent Vector Predictor	0.0141	0.0014	0.0177	0.1778
Full-trend Predictor	0.0095	0.0005	0.0122	0.1101

The full-trend model achieved significantly lower DTW, indicating better alignment with the temporal dynamics of the plant parameters during transient events. This performance can be attributed to the use of the decoder during training, which enforces learning of the end-to-end mapping between latent space and observable signals. By optimizing prediction in signal space, the model learns to focus on patterns that matter most for accurate trajectory forecasting.

A qualitative analysis was performed on a LOCA scenario with a stuck-open pressurizer relief valve (PORV), focusing on three key variables: Cold Leg Temperatures (Loops 2 & 3) and Containment Pressure. **Figure 2** shows the prediction results compared to actual values, with a shaded band indicating a ±5% error range. The full-trend model's output closely follows the actual signal in all three variables, capturing both transient and steady-state behaviors more effectively than the latent-vector model.

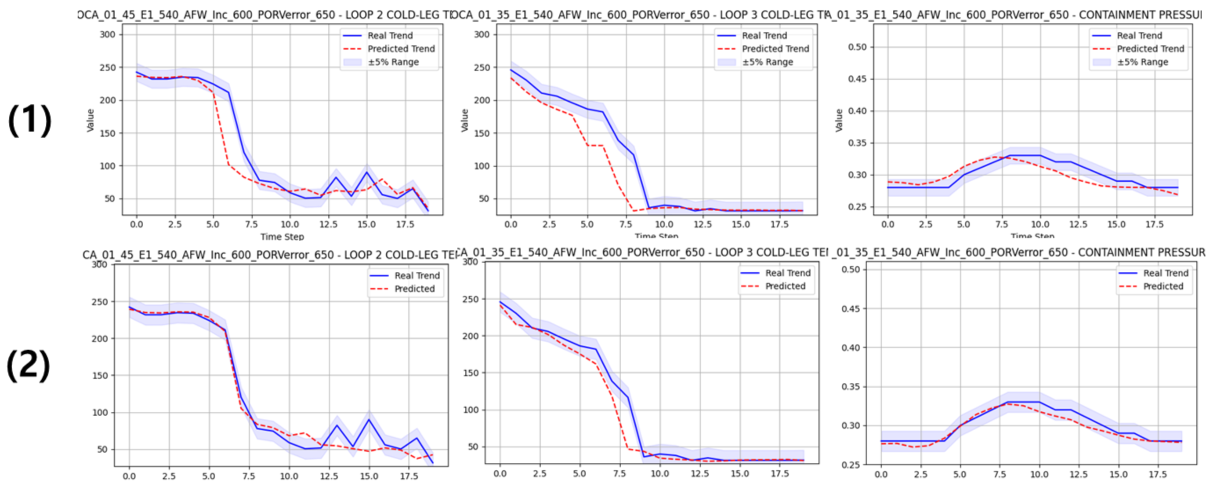

Figure 2 Trend prediction results for three key parameters: (1) shows the prediction results using latent vector loss, while (2) shows results using the full-trend loss strategy with a frozen decoder.

This study presents a latent-space-based trend prediction framework for real-time operator support in small modular reactors. The combination of a deep LSTM Autoencoder and a decoder-assisted predictor results in significantly improved accuracy and temporal alignment in accident scenarios. Future work will explore validation strategies for untrained scenarios, as well as integrate uncertainty quantification methods to improve the robustness and interpretability of predictions in digital control room environments.

ACKNOWLEDGMENTS

This work was supported by Korea Institute of Energy Technology Evaluation and Planning(KETEP) grant funded by the Korea government(MOTIE)(RS-2024-00403194, Next-Generation Nuclear Technology Creation IP-R&D Talent (Human Resources) Development Project)

This research was supported by the National Research Council of Science & Technology(NST) grant by the Korea government (MSIT) (No. GTL24031-400)

REFERENCES

[1] Bae, J., Kim, G., & Lee, S. J. (2021). Real-time prediction of nuclear power plant parameter trends following operator actions. *Expert Systems with Applications*, 186, 115848.

[2] Kim, H., & Kim, J. (2023). Long-term prediction of safety parameters with uncertainty estimation in emergency situations at nuclear power plants. *Nuclear Engineering and Technology*, 55(5), 1630–1643.

[3] Bae, J., Park, J.W., & Lee, S.J. (2022). Limit surface/states searching algorithm with a deep neural network and Monte Carlo dropout for nuclear power plant safety assessment. *Applied Soft Computing*, 124, 109007.