Data-driven reduced order models for power reconstruction during load-follow operation

Dongmin Yun, Deokjung Lee*
Department of Nuclear Engineering, Ulsan National Institute of Science and Technology, UNIST-gil 50, Eonyang-eup, Ulju-gun, Ulsan, 44919
**Corresponding author: deokjung@unist.ac.kr*

Keywords : load-follow operation, proper orthogonal decomposition, dynamic mode decomposition, sparse identification of nonlinear dynamics

1. Introduction

The operation of a single reactor unit is increasingly characterized by flexible operation to accommodate the intermittent solar generation profiles which depend on the existence of sunlight. Under random load-following operation (LFO), the real-time dynamics of control rod movement, combined with spatiotemporally evolving fission product behavior and resulting dynamic power imbalance, pose a fundamental challenge: existing power distribution synthesis methodologies lack the adaptive capability to track actual core dynamics under stochastic operating conditions [1].

Reduced order models (ROMs) offer a data-driven alternative to bypass such computationally intensive full-order simulations [2]. By learning from high-fidelity 3-D nodal diffusion simulations of load-follow transients, ROMs can compress complex spatial-temporal dynamics into low-dimensional manifolds and extrapolate to unseen operating points without offline computation, enabling real-time prediction under arbitrary control policies and setpoint trajectories [3][4].

In this study, the high-fidelity and pseudo-realistic simulation data of load-following operation are generated by STREAM/RAST-K [5], which has been developed in UNIST CORE Lab. As a dimensionality reduction method, the proper orthogonal decomposition (POD) has been implemented using QR-based singular value decomposition (SVD) [6]. This POD basis is then used for both projection and reconstruction. Two complementary data-driven methods are employed to model the temporal evolution in latent space: Dynamic Mode Decomposition with control (DMDc) [7] which extracts a linear affine model based on Koopman operator theory, and Sparse Identification of Nonlinear Dynamics with control (SINDy-c) [8] which discovers sparse nonlinear models into first order differential equation. The aim of this study is to reduce the generalization error of data-driven power reconstruction methods, comparing two latent propagation methods, DMDc and SINDy-c.

2. Methods and Results

In this section, the whole procedure of data-driven ROMs for power reconstruction from data generation to inference stage are described. The schematic flowchart of this method is illustrated in Fig. 1.

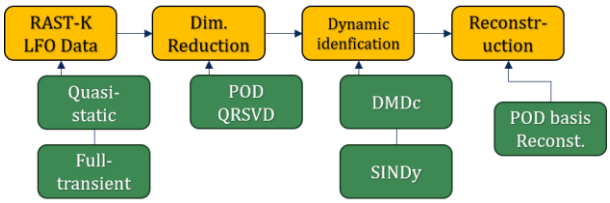


Fig. 1. Study overview of data-driven ROM for power reconstruction

2.1 Data generation : RAST-K

To generate a comprehensive dataset, we use a mutually complementary two-stage simulation strategy that iterates the full-transient and quasi-static approaches. We integrate the two simulation stages to balance high-fidelity transient representation with broad parametric coverage. The equilibrium cycle of APR1400 is used for data generation of this study.

First, we implemented full-transient simulations using the RAST-K's pseudo module of reactivity regulation system (RRS) [9], which automatically maneuver the control rod and boron concentration to follow target core power and reference axial shape index (ASI) in real-time. From these, the detailed control histories of various 2-day LFO scenario cases on the beginning of cycle (BOC) have been captured.

Second, using this history, quasi-static simulations of boron search mode in whole cycle of such randomized schedule have been calculated as a prototype of dataset. Then, restarting from LFO results, the extracted control history over middle of cycle (MOC) and end of cycle (EOC) has been iteratively captured with long term LFO effects. Each LFO scenario consists of 50%, 70%, and 80% of target load-following and temporal patterns of 3-6-3-12, 2-8-2-12, and 2-6-2-14.

Finally, the whole-cycle quasi-static using the chronicle LFO control history are calculated with 15-minute time stepping then the data are acquired using HDF5 format to avoid truncation issue of ascii-style edits.

2.2 Data preprocessing

The dataset represents a three-dimensional state space comprising assembly-wise power and burnup, and the number density distribution of xenon-135 and samarium-149. This spatial discretization of assembly geometry and axial extent yields thousands of flattened spatial nodes per physical variable X ($\in \mathbb{R}^{n\times m}$), where n is the

number of time-step and m is the number of node-wise features. We represent control actions by the nine control rod bank positions, the primary-coolant boron concentration, and the power setpoint; these inputs form the control data $u(\in \mathbb{R}^{n\times p})$. The min-max normalization has been applied feature-wise. The validation dataset has totally different LFO schedule and control history from train dataset to verify the generalization error of this study. The main control history set of one case from train dataset are illustrated in Fig. 2.

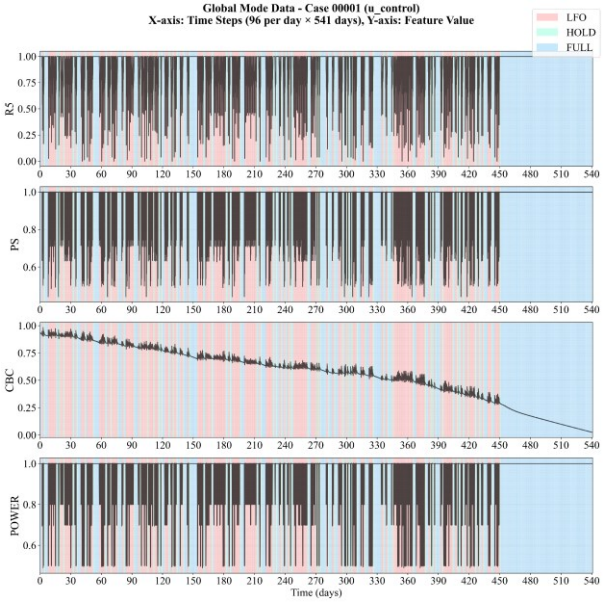


Fig. 2. The normalized control data (u) of sampled case

2.3 Dimensionality reduction : POD

POD serves as the foundation for reduced-order modeling by extracting dominant spatial-temporal structures from high-dimensional system observations. In this work, POD projects the full-dimensional reactor power distribution into a low-dimensional latent representation that preserves essential system dynamics while achieving substantial computational savings. The decomposition is computed using QR-factorization-based singular value decomposition (QRSVD) implemented in a distributed data-parallel (DDP) configuration across 8 graphic processing units (GPU), enabling efficient processing of large-scale snapshot ensembles.

The SVD is then computed on the much smaller QR-decomposed matrix R ($X = QR$):

$$R = \tilde{U}\Sigma V^T \tag{1}$$

The full left singular vectors of the original matrix are recovered via:

$$\boldsymbol{U} = Q\tilde{U} \tag{2}$$

the SVD computation operates on an $R \in \mathbb{R}^{m\times m}$ matrix rather than the much larger $X \in \mathbb{R}^{n\times m}$ matrix, reducing computational cost from $\mathcal{O}(n^2m)$ to $\mathcal{O}(m^3)$.

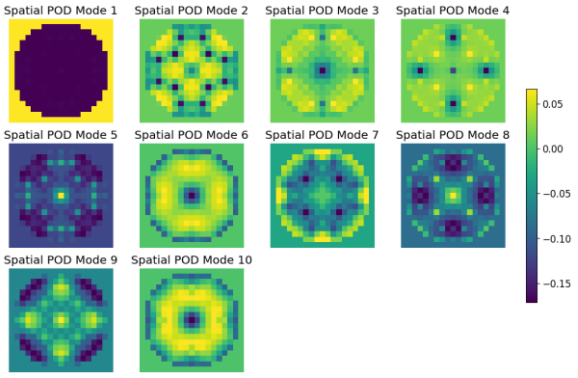


Fig. 3. Top 10 ranked POD basis (U_r)

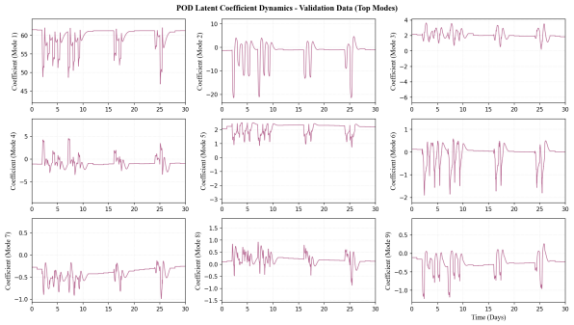


Fig. 4. Top 9 ranked latent coefficients $(a_r = \Sigma_r V_r^T)$

Fig. 3 shows the POD results of high-ranked principle spatial components, especially marked at the control rod positions, and Fig. 4 describes the latent dynamics including the temporal components of the system.

The retained POD rank r is selected by "energy" criterion of $E(r) = \sum_{i=1}^{r}\sigma_i^2 / \sum_{i=1}^{m}\sigma_i^2 \geq 99\%$, where the singular values $\sigma \in \Sigma$ represent the relative importance of each mode. The reconstruction is from POD basis from train dataset $(U_{r,train})$ and latent coefficients from validation dataset $(a_{r,valia})$. The performance metric is the mean relative error (MRE) between true power distribution from validation dataset (X_{valid}) and reconstructed power distribution ($X_{recon} = U_{r,train} \cdot a_{r,valia}$). Fig. 5 illustrates MRE of POD methods using the validation dataset. The MRE is mostly under 0.1% which means the linearly projected manifold for reduction of dimensions is well generalized enough.

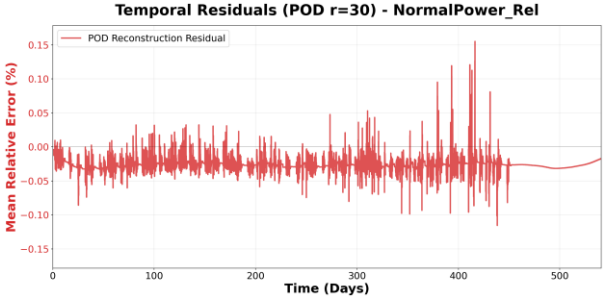


Fig. 5. MRE(%) of reconstruction of latent variables on

validation dataset using POD xenon basis from trained dataset

2.4 Linear-affine Koopman operator: DMDc

Koopman operator $\mathcal{K}$ is an infinite-dimensional linear operator acting on observables of a dynamical system. For a discrete-time control system of the form $\boldsymbol{a}_{k+1} = f(\boldsymbol{a}_k, \boldsymbol{u}_k)$ the Koopman operator advances any observable g by one time step:

$$\mathcal{K}g(\boldsymbol{x}_k) = g(\boldsymbol{x}_{k+1}) = g\big(f(\boldsymbol{x}_k, \boldsymbol{u}_k)\big) \tag{3}$$

In the context of POD-based model reduction, we select the latent coordinates vector $\boldsymbol{a}_k \in \mathbb{R}^r$ (resulting from POD projection) as those observables. This choice leads to a control-affine linear model in the latent space:

$$\boldsymbol{a}_{k+1} = \boldsymbol{A}\boldsymbol{a}_k + \boldsymbol{B}\boldsymbol{u}_k \tag{4}$$

, where $\boldsymbol{A} \in \mathbb{R}^{r\times r}$ is the identified Koopman operator matrix in the latent space, and $\boldsymbol{B} \in \mathbb{R}^{r\times p}$ is the control input matrix; r denotes the number of retained POD ranks, and p is the dimension of the control input.

To identify the matrices $\boldsymbol{A}$ and $\boldsymbol{B}$, we construct snapshot matrices from the available training data. Let $\boldsymbol{a}^{(l)} \in \mathbb{R}^{r\times(n-1)}$ denote the matrix of latent snapshots and let $\boldsymbol{u}^{(l)} \in \mathbb{R}^{p\times(n-1)}$ denote the matrix of control inputs at times $k = l, l+1, \dots, l+n-1$. The least-squares identification problem seeks to minimize the following Frobenius norm:

$$\min_{\boldsymbol{A},\boldsymbol{B}} \left\| \boldsymbol{a}^{(1)} - \boldsymbol{A}\boldsymbol{a}^{(0)} - \boldsymbol{B}\boldsymbol{u}^{(0)} \right\|_F^2 \tag{5}$$

The resulting control-affine model is then validated through multi-step predictions on held-out test data, in Section 2.6.

2.5 Sparse Identification of Nonlinear Dynamics with control (SINDy-c)

While POD and Koopman operator theory provide an efficient framework for real-time power reconstruction, the assumption of linear dynamics introduces systematic errors under transient conditions and nonlinear feedback regimes characteristic of load-follow operations. To address these limitations while maintaining computational tractability, this section introduces SINDy-c, a methodology that identifies interpretable, explicit governing equations of the form $\frac{d\boldsymbol{a}}{dt} = \boldsymbol{f}(\boldsymbol{a}, \boldsymbol{u}) = \Theta(\boldsymbol{a}, \boldsymbol{u})\boldsymbol{\Xi}$ as sparse linear combinations of interpretable basis functions (Θ) – including polynomial, exponential, Fourier and control-state coupling terms – and corresponding sparse coefficients ($c \in \boldsymbol{\Xi}$). This method provides stronger robustness to be a generalized model, using a concept of 'sparsity'.

An optimization problem in finding a sparse vector $\boldsymbol{\Xi}$ that fits the derivative $\dot{\boldsymbol{a}}$ is performed least squares fitting with L1 regularization using λ as parameter:

$$c_k = \underset{c_k'}{argmin}\, ||\dot{\boldsymbol{a}}_k - \Theta(\boldsymbol{a}, \boldsymbol{u})c_k'||_2 + \lambda||c_k'||_1 \tag{6}$$

The sequentially thresholded least square (STLSQ) algorithm combining ridge regression, thresholding, and iterative refinement to isolate physically dominant dynamics while suppressing minor terms with given sparsity. Fig. 6 illustrates the heatmap of highlighted coefficient and the representative terms of trained SINDy-c equation.

- STLSQ algorithm

1) Ridge Regression (initialize all terms)

$$\boldsymbol{c}^* = (\boldsymbol{\Theta}^T\boldsymbol{\Theta} + \alpha\boldsymbol{I})^{-1}\boldsymbol{\Theta}^T\frac{\boldsymbol{da}}{\boldsymbol{dt}}$$

2) Thresholding (enforce sparsity)

$$\boldsymbol{c}_j^{thresh} = \begin{cases} \boldsymbol{c}_j^* & if\ |\boldsymbol{c}_j^*| \geq \boldsymbol{\tau} \\ \boldsymbol{0} & otherwise \end{cases}$$

3) Iterative Refinement (repeat 1-2 until converged)

$$\boldsymbol{c}^{(k+1)} = (\boldsymbol{\Theta}_{active}^T\boldsymbol{\Theta}_{active} + \alpha\boldsymbol{I})^{-1}\boldsymbol{\Theta}_{active}^T\frac{\boldsymbol{da}}{\boldsymbol{dt}}$$

4) Return the converged $c \in \boldsymbol{\Xi}$ (final)

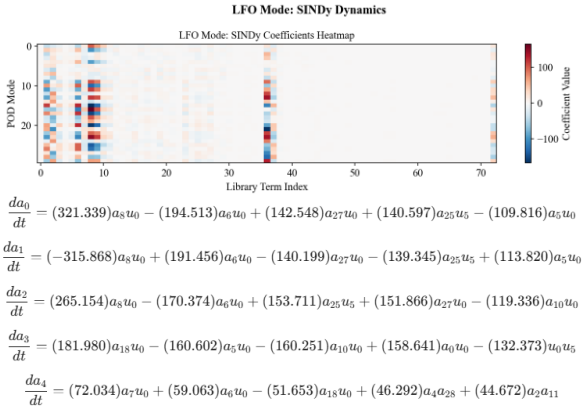


Fig. 6. Trained SINDy-c model example (top 5 terms are sampled for visualization)

2.6 Comparison for DMDc and SINDy-c

The validation results of DMDc and SINDy-c model are figured. Fig. 7 and Fig. 9 show the prediction of top 5 ranked latent coefficients and those errors from validation dataset. Fig. 8 and Fig. 10 describe the mean absolute relative errors (MARE) of the reconstructed power distribution at timestep t, denoted in Eq. (7), which are predicted by each dynamic identification method, by using POD basis which is preset from train dataset. Fig. 7 and 8 are tested on day 0 to 9, and Fig. 9 and 10 are tested on day 400 to 409.

$$\mathrm{MARE}(t) = \frac{1}{n_x n_y n_z}\sum_{p=1}^{n_x n_y n_z}\left|\frac{X_{p,t}^{recon} - X_{p,t}^{val}}{X_{p,t}^{val}}\right| \tag{7}$$

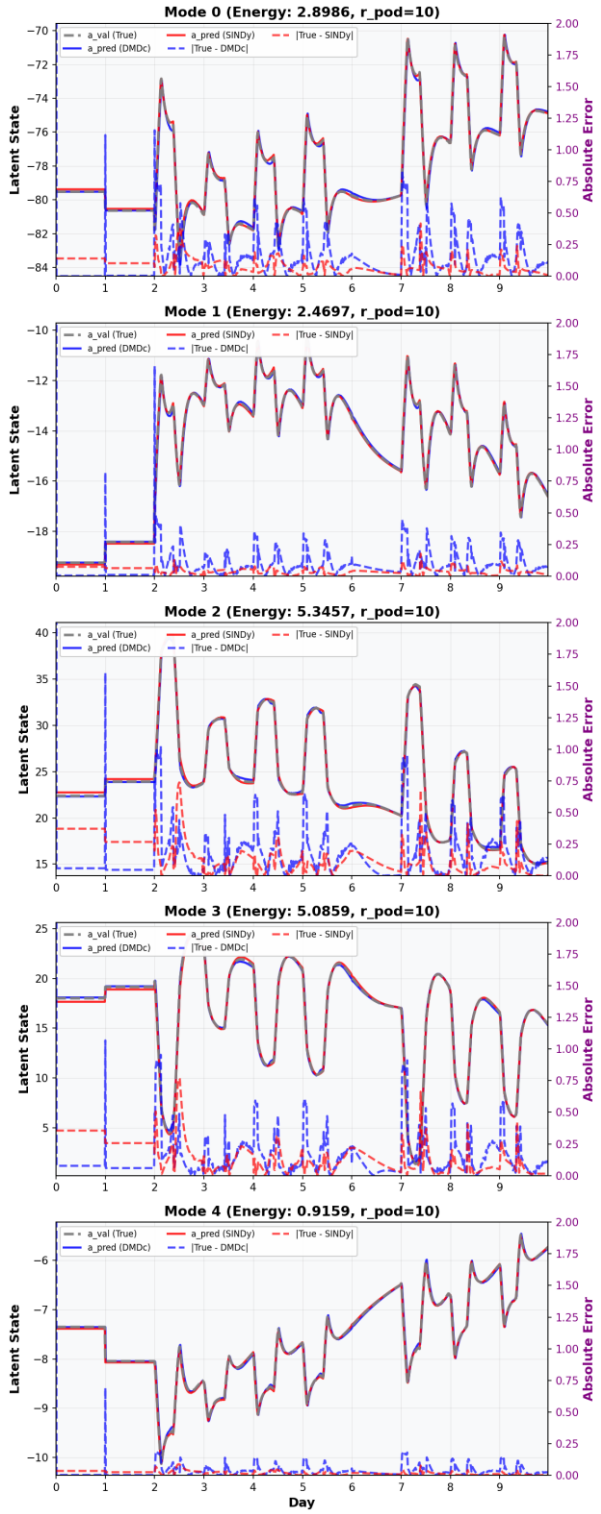


Fig. 7. One-step predictions and errors of latent state coefficients at day 0 to 9

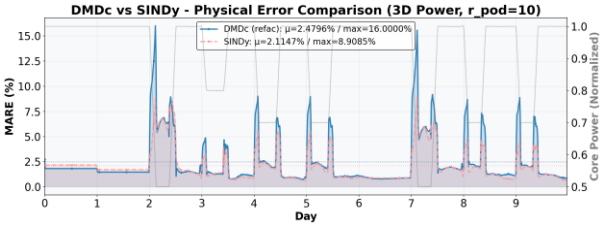


Fig. 8. MARE of 3-D power reconstruction at day 0 to 9

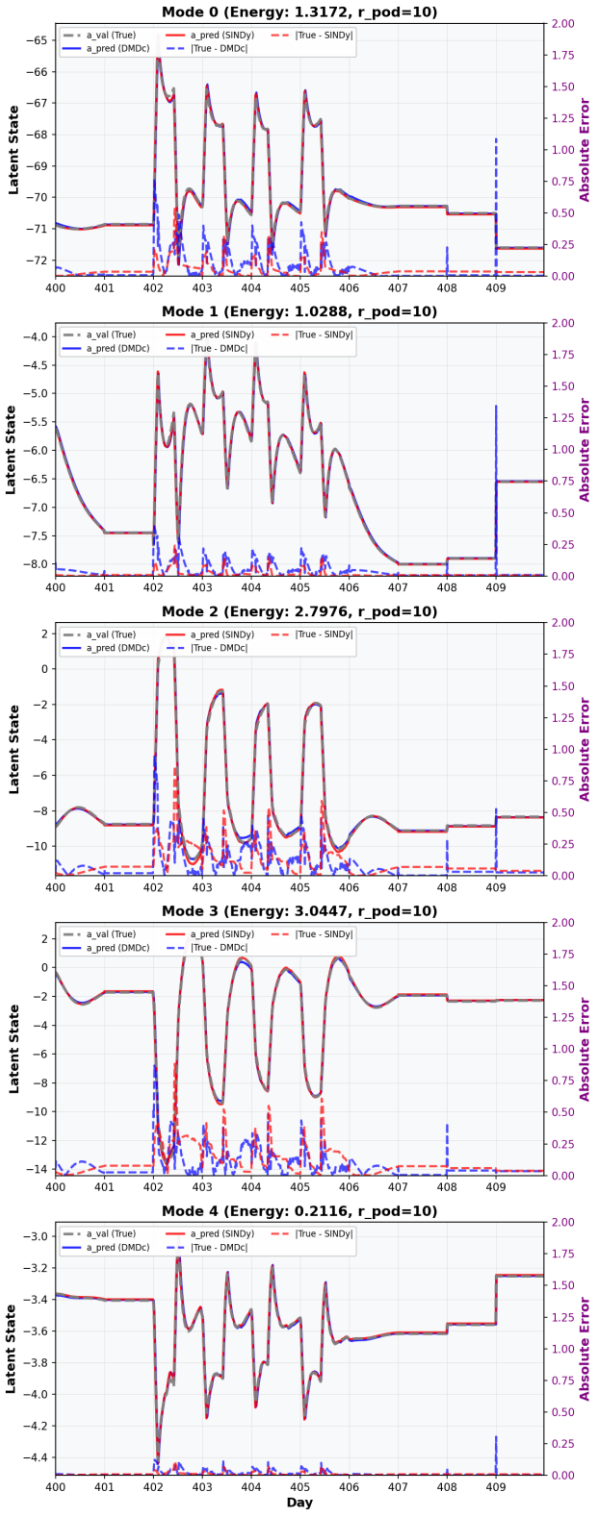


Fig. 9. One-step predictions and errors of latent state coefficients at day 400 to 409

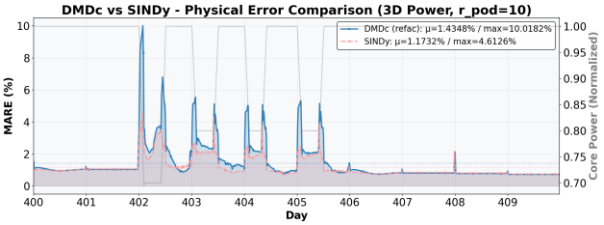


Fig. 10. MARE of 3-D power reconstruction at day 400 to 409

From results, both DMDc and SINDy-c exhibit higher reconstruction MAREs during load-follow operation than during non-load-follow periods, with the magnitude of the reconstruction error generally increasing with the extent of the power maneuver. In the validation case, load-follow operation was performed on 7 out of the first 10 days of the cycle. During this period, DMDc yielded an average MARE of 2.48% and a maximum MARE of 16.00%, whereas SINDy-c achieved an average MARE of 2.11% and a maximum MARE of 8.91%. Although the average MAREs of the two dynamic identification methods are comparable; a notable difference is observed during transitions between different daily load-following scenarios.

In particular, DMDc exhibits reconstruction errors approximately twice as large as those of SINDy-c during power reduction and recovery transients. This behavior can be attributed to the relatively large latent-state errors occurring at the transition points between different daily load-following scenarios, as well as the limitations of the linear-affine DMDc formulation in representing the nonlinear behavior associated with control-rod movements and xenon dynamics. In contrast, the nonlinear sparse representation provided by SINDy-c allows these transient dynamics to be represented more effectively, resulting in lower reconstruction errors during changes in reactor power.

3. Conclusions

This work demonstrates a data-driven reduced-order modeling (ROM) framework for three-dimensional (3-D) power distribution reconstruction during load-follow operation in nuclear reactors. The proposed framework integrates full-cycle load-follow operation simulations using RAST-K, dimensionality reduction based on QR-based POD, and latent-space dynamics identification using DMDc and SINDy-c. The results show that the ROM efficiently extracts the dominant spatial–temporal features of the reactor core while maintaining numerical stability through orthogonal transformations, yielding a low-rank latent representation suitable for dynamic model identification and real-time reconstruction.

The proposed framework can be extended beyond 3-D power distribution reconstruction to the prediction of other high-dimensional, non-measurable core states, such as the spatial distributions of xenon and burnup. This capability provides a potential means of predicting the highly intermittent and dynamically varying behavior of reactors operating under flexible load-follow conditions in a more comprehensive and generalizable manner.

The proposed framework was trained exclusively using RAST-K simulations. Therefore, discrepancies between simulated core physics and actual reactor behavior need to be characterized using commissioning and operational data. In addition, the current evaluation was performed using simulated validation data only. For application to actual load-follow operation, real-time SPND in-core instrumentation signals should be incorporated to provide online measurement feedback for latent-state correction. Furthermore, to quantitatively assess the improvement in the generalization capability of the ROM, power reconstruction using a full-system data-driven approach should be evaluated as an independent verification method.

ACKNOWLEDGEMENTS

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. RS-2025-16063033).

REFERENCES

[1] H. Yang, Z. Li, S. Chen, T. Wang, H. Wang, and S. Chen, Review of methods for online power-distribution reconstruction, Annals of Nuclear Energy, Vol.221, 2025.
[2] N. Abrate, S. Dulla, and N. Pedroni, A non-intrusive reduced order model for the characterisation of the spatial power distribution in large thermal reactors, Annals of Nuclear Energy, Vol.184, 2023.
[3] B. Peherstorfer, and K. Wilcox, Data-driven operator inference for nonintrusive projection-based model reduction, Computer Methods in Applied Mechanics and Engineering, Vol.306, p.196, 2016.
[4] A. Mamonov, and M. Olshanskii, Tensorial parametric model order reduction of nonlinear dynamical systems, SIAM Journal of Scientific Computing, Vol.46, p.A1850-A1878,2024.
[5] J. Park, J. Jang, H. Kim, J. Choe, D. Yun, P. Zhang, A. Cherezov, and D. Lee, RAST-K v2- Three Dimensional Nodal Diffusion Code for Pressurized Water Reactor Core Analysis, Energies, Vol.13, p. 6324, 2020.
[6] J. Demmel, L. Grigori, M. Hoemmen, and J. Langou, Communication-optimal parallel and sequential QR and LU factorizations: theory and practice, SIAM Journal of Scientific Computing, Vol.34, 2012.
[7] J. Tu, C. Rowley, D. Luchtenburg, S. Brunton, and J. Kutz, On dynamic mode decomposition: theory and applications, Journal of Computational Dynamics, Vol.1, 2014.
[8] S. Brunton, J. Proctor, and J. Kutz, Discovering governing equations from data by sparse identification of nonlinear dynamical systems, Proc. Natl. Acad. Sci. U.S.A., Vol.113, p.3932-3937, 2016.
[9] D. Yun, H. Sohn, and D. Lee, Daily load-follow simulation using 3-D transient and quasi-static solver of RAST-K, Proceedings of the Reactor Physics Asia 2025 (RPHA2025) Conference, 2025.