INL/CON-25-82713-Revision-0

Physics-Informed Recurrent Neural Networks to Predict Reactor Operations of the AGN-201 Nuclear Reactor

April 2025

Jaden Sonny Palmer, Ryan Hunter Stewart, Ashley Jeanette Brockman Shields

Changing the World's Energy Future

INL is a U.S. Department of Energy National Laboratory operated by Battelle Energy Alliance, LLC

INL/CON-25-82713-Revision-0

Physics-Informed Recurrent Neural Networks to Predict Reactor Operations of the AGN-201 Nuclear Reactor

Jaden Sonny Palmer, Ryan Hunter Stewart, Ashley Jeanette Brockman Shields

April 2025

Idaho National Laboratory
Idaho Falls, Idaho 83415

http://www.inl.gov

Prepared for the
U.S. Department of Energy
Under DOE Idaho Operations Office
Contract DE-AC07-05ID14517

Recurrent Neural Networks to Predict Reactor Operations using Physics-Informed Data and Rolling Statistics

Jaden Palmer[1], Ryan Stewart[2], Ashley Shields[2]

[1]Idaho State University, Pocatello, Idaho
[2]Idaho National Laboratory, Idaho Falls, Idaho
palmjad2@isu.edu

INTRODUCTION

To address the growing concerns of climate change, advanced nuclear reactors serve as a possible solution to achieve carbon neutrality goals. The International Atomic Energy Agency (IAEA), which is tasked with providing assurance of peaceful uses of nuclear material, has seen little growth in resources compared to the unprecedented expansion of the nuclear industry [1]. New proliferation detection methods are necessary to accurately monitor nuclear states to ensure compliance with the Nuclear Non-Proliferation Treaty set by the IAEA [2]. Digital twins (DTs), defined by the Digital Twin Consortium as "an integrated data-driven virtual representation of real-world entities and processes, with synchronized interaction at a specified frequency and fidelity," serve as a possible solution [3]. The current state-of-the-art DT for the AGN-201 reactor is comprised of multiple models, particularly physics and machine learning (ML) models [4, 5]. These models work in synchrony to accurately predict transient reactor behavior during an operation and monitor for anomalous behavior occurs.

When modeling the operational capabilities of the newest generation of reactors (Generation IV), many have no current operating plants. Therefore, it will be imperative that modeling can capture expected operations to be adapted once real operations begin. Training an ML model on physics-informed synthetic data serves as a possible solution to modeling reactors prior to their construction. ML models, specifically in the form of physics-informed neural networks, have already been used to predict reactor transients by solving point kinetics equations with temperature feedback [6]. This work aims at predicting reactor operations of Idaho State University's (ISU's) AGN-201 reactor, using a long short-term memory (LSTM) recurrent neural network (RNN) trained on physics-informed synthetic data. Predictions made by the LSTM were validated on historical operational data from the AGN-201 reactor.

AGN-201 Reactor

The nuclear testbed utilized for this research is ISU's 5 watt AGN-201 teaching and research reactor [7]. A top view of the AGN-201 is shown in Figure 1. The reactor power is measured by three channels, shown in the diagram as Channels 1, 2, and 3, which measure the startup count rate, power on a logarithmic scale, and power on a linear scale respectively. Channels 2 and 3 measure the reactor power during an operation, while Channel 1 is specifically used during reactor startup. The measurements made by Channels 2 and 3 are similar but not identical as the nature of the scales that both measure power allow for different scram conditions. Channel 2 was designated as the target variable predicted by the ML model.

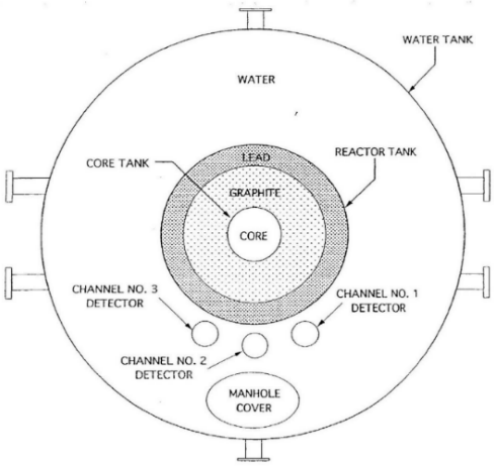

Fig. 1. AGN-201 top view.

The fuel for the AGN-201 is stored inside of its control rods, which are inserted from the bottom of the reactor up to a full insertion height of 25 cm. During an operation, two safety control rods are always fully inserted into the reactor, and a coarse control rod (CCR) is also generally at or near full insertion. Each of these control rods have a total reactivity value of \$1.68. The fine control rod (FCR) serves as a smaller control rod with a worth of \$0.42 and is nominally used to manipulate the reactor's power.

The temperature of the water shielding surrounding the reactor core is also measured. However, due to the smaller scale of the AGN-201 compared to conventional reactors, the temperature of the water shielding is generally measured to be the ambient room temperature of the reactor bay as the reactor core does not create an appreciable heat flux. Therefore, the reactor power at which the AGN-201 is operating has no correlation with the temperature of the water

reflector surrounding the reactor core. As the temperature measurements, which are generally in correspondence to the power produced by larger reactors, are insignificant for accurate ML predictions of the AGN-201, synthetic Channel 3 data were initially used to supplement the model's training data.

METHODOLOGY

Synthetic Data Generation

Using material and geometric data taken from the safety analysis report and from previous benchmarking of the AGN-201, a Serpent Monte Carlo model of the reactor was created [8, 9, 10]. Using the reactor's characteristics detailed from these reports, the CCR and FCR reactivity worths from the model were calculated to be \$1.86 and \$0.52. This difference from the original control rod worths compared to physical reality resulted in a 200 per cent mille (pcm) bias between the known critical value and the Serpent model. This bias was accounted for by allowing the reactor to be at a critical state when k_{eff} was 1.00225, allowing for 25 pcm of uncertainty in the Serpent model's approximate bias.

Parameters predicted by the Serpent model, such as the delayed neutron fractions and decay rates, were used to train a point kinetics equations surrogate model (PKE-SM) to create synthetic reactivity and power data. The PKE-SM was created via the *openpointkinetics* Python module [11]. Using the constants provided by Serpent, the PKE-SM solves for reactivity through Equation 1 and the change in neutron density over time (a parameter correlated to power) through the coupled seven ordinary differential equations presented in Equations 2 and 3, where "i" designates six groups of delayed neutron precursors.

$$\rho(t) = \frac{k_2 - k_1}{k_2 \cdot k_1} \quad (1)$$

$$\frac{dn}{dt} = \left[\frac{\rho(t) - \beta}{\Lambda}\right] n(t) + \sum_{i=1}^{6} \lambda_i C_i(t) \quad (2)$$

$$\frac{dC_i}{dt} = \frac{\beta_i}{\Lambda} n(t) - \lambda_i C_i(t) \qquad i = 1, \ldots, 6 \quad (3)$$

Due to the internal physics of the Channel 2 and Channel 3 detectors, mechanical vibrations, and possible electrical interference, there is an uncertainty present in power measurements, conventionally characterized as "noise." To ensure that synthetic data represented this uncertainty, artificial noise was added to the synthetic data. The percent standard deviation (PSD) was calculated to quantify the magnitude of noise for the Channel 2 and Channel 3 detectors. Artificial noise was added using a random normal distribution, using the calculated PSD. Measurements made below the threshold of 15 mW of power were also found to result in a higher magnitude of noise; therefore, two sets of PSDs were calculated for each sensor. A PSD of 1.72% and 0.293% was found for Channels 2 and 3, respectively, above 15 mW, while below 15 mW, the PSDs were 10.48% and 7.14%, respectively. The resulting effect of adding this artificial noise onto a synthetic dataset is illustrated in Figure 2. In total, 180 synthetic datasets with different control rod movements and initial power conditions were created using this methodology to train the LSTM.

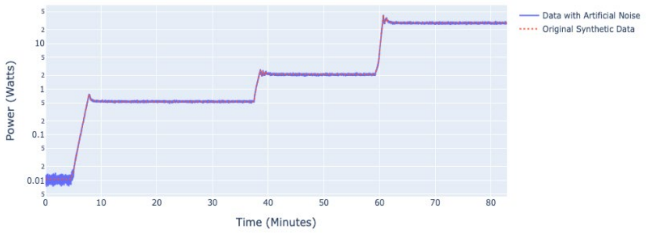

Fig. 2. Noisy vs. initial synthetic data from PKE-SM.

Long Short-Term Memory Recurrent Neural Network

The PKE-SM informed by Serpent may accurately predict reactor operations but fails to account for sensor drift over time, which is a physical reality for reactors [12]. Implementing an ML model can capture changes in sensor behavior and alert users when retraining of the surrogate models is necessary. An LSTM RNN was trained to account for sensor drift in predicting operations of different reactor types using different feature variables.

The LSTM model was chosen because of the architecture's strength in predicting time-series data by understanding short- and long-term patterns within the data [13]. Two layers consisting of 64 and 32 LSTM units were used. Each LSTM unit contains a cell state, an input gate, an output gate, and a forget gate. The cell state is present throughout each gate, serving as the unit's internal memory. The forget gate manages the removal of information from a cell state, while the input gate decides which new information from the current input should be incorporated into the cell state. The output gate determines the next hidden state by adjusting the model weights, reflecting the information at the current time step.

Validation Data Preprocessing

This work focused on operations after the reactor was brought to a critical state (i.e., k_{eff} = 1). Operational data retrieved from the AGN-201's data acquisition system to validate the PKE-SM and LSTM were found to contain periods of noncritical data prior to startup and following shutdown. To ensure accurate validation of virtual models, an autonomous critical data extrapolator algorithm was developed for efficient preprocessing of validation data. The algorithm identifies critical data by detecting reactor startup and shutdown events and eliminates data that occur before and after these events. Reactor startup is recognized by an initial supercritical state with a high neutron production rate,

followed by stabilization at a k_{eff} of 1. Shutdown is detected by a rapid drop in power and control rod positions near zero centimeters. This approach ensures that only critical data, where the reactor operates stably, are used for validating the PKE-SM and LSTM.

Feature Engineering

Initial results from the LSTM training utilize Channel 3 data in the absence of meaningful temperature data. To ensure the LSTM can create accurate predictions without relying on power sensor data, methods to eliminate Channel 3 feature data from the training dataset are currently under investigation. Without a measured heat flux and Channel 3 data, the initial power will be unknown to the LSTM. With this understanding, power data is normalized to the first datapoint of the target dataset, allowing the LSTM to learn power manipulations from reactor's initial condition. Additionally, two feature variables related to the dynamics of the reactor were added to the training dataset, reactivity and the inverse reactor period. Reactivity was calculated by the PKE-SM and is determined by the Channel 2 and Channel 3 sensors. Inverse period is also computed by the Channel 2 sensor and calculated through Equation 4. The prompt neutron lifetime, l_p, was found to be 4.160E-05 seconds for the AGN-201 reactor from the Serpent model, similar to the prompt neutron lifetime of 4.283E-5 seconds found in [14].

$$\frac{1}{T} = \frac{k_2 - k_1}{l_p} \quad (4)$$

In addition to the added reactor physics features, rolling statistics were also implemented into the total training dataset. A rolling average, average ratio, standard deviation, maximum, median, skew, and kurtosis for CCR, FCR, inverse period, and reactivity feature variables were implemented for 3, 30, and 300 second rolling periods, as these periods are multiples of the natural reactor period. The summary statistics, such as kurtosis and skew which are presented by Equations 5 and 6, respectively, over the time windows provide added insight into changes in distributional shape for those time periods. Adding rolling statistics to the feature dataset capitalizes on the LSTM's strength in interpreting temporal data compared to other ML architectures, providing the model a more comprehensive understanding compared to raw observations alone.

$$\text{Skew}_{\text{rolling}} = \frac{1}{N}\sum_{t=0}^{N-1}\left(\frac{x_t - \mu}{\sigma}\right)^3 \quad (5)$$

$$\text{Kurt}_{\text{rolling}} = \frac{\sum_{i=0}^{N-1}(x_{t-i} - \mu)^4 / N}{\sigma^4} \quad (6)$$

RESULTS AND ANALYSIS

Figure 3 displays PKE-SM predictions for reactivity compared to physical data taken from the AGN-201 on February 10, 2023. Initial PKE-SM predictions, designated by PKE-SM-CF1, tend to overestimate reactivity measurements made by the Channel 2 sensor during transient periods. This overestimation is further propagated when the PKE-SM solves for power (as reactivity is integrated to solve Equations 2 and 3), as shown in Figure 4. This overprediction originates from the 200 pcm bias in the Serpent model where the reactor constants used to solve the PKEs are derived. To account for this bias, a reactivity correction factor (RCF) was implemented to scale Equation 1. Generally, an RCF between 0.75 and 0.80 yielded the best results in synthetic data generated by the PKE-SM compared to operational datasets. This RCF plays a crucial role in ensuring that the magnitude of power during an operation is accurately predicted and monitored by the PKE-SM and DT, respectively.

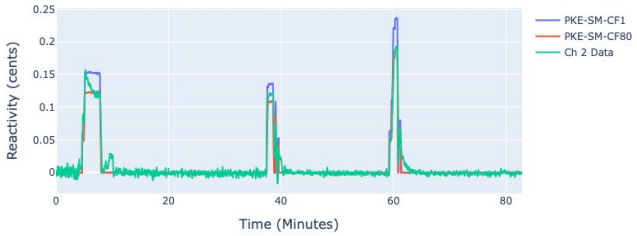

Fig 3. Comparison of the PKE-SM reactivity predictions.

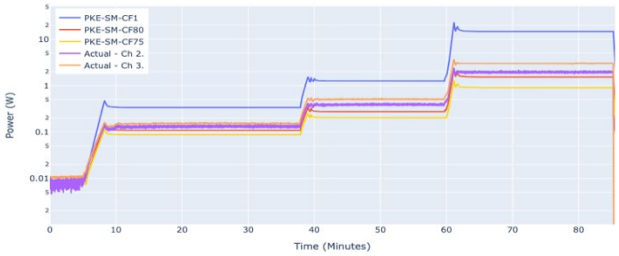

Fig 4. Comparison of the PKE-SM power predictions.

LSTM predictions compared to the operational dataset are shown in Figure 5 for various RCFs. Table 1 presents the accuracy of these predictions through the R^2, mean absolute error (MAE), and mean absolute percent error (MAPE) metrics. LSTM models showed minimal difference in performance with or without correction factors. Analysis revealed that the synthetically created feature data, designed to mimic Channel 3 measurements, heavily influenced the prediction accuracy of Channel 2.

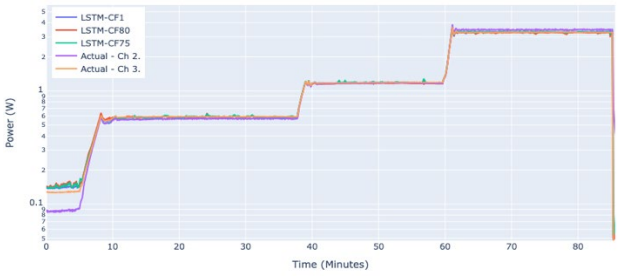

Fig 5. Comparison of LSTM-CF predictions.

Table 1. LSTM accuracy for various correction factors

Correction Factor	R^2	MAE (W)	MAPE
1.0	0.984	0.095	4.4%
0.80	0.986	0.082	4.3%
0.75	0.986	0.078	4.11%

CONCLUSIONS AND FUTURE WORK

An LSTM RNN trained on physics-informed synthetic data effectively predicted operations for ISU's AGN-201 reactor. To validate LSTM predictions and PKE-SM calculations, an autonomous critical data extrapolator algorithm was developed to preprocess operational data, ensuring high fidelity by eliminating noncritical periods. The model showed minimal performance differences with or without RCFs since the synthetic data mimics Channel 3 measurements, which significantly influences Channel 2 predictions. Future work involves implementing proposed feature engineering methods to reduce the LSTM's reliance on the Channel 3 feature variable. Implementing these new reactor physics and rolling statistics feature variables will increase the importance of the RCF in training ML models for predicting reactor operations in the absence of power data.

NOMENCLATURE

k_{eff} = Effective Neutron Multiplication Factor
k_1 = Critical Effective Neutron Multiplication Factor
k_2 = Transient Measured/Calculated Effective Neutron Multiplication Factor
$n(t)$ = Neutron Population
$\rho(t)$ = Reactivity
β = Delayed Neutron Fraction
λ = Decay Constant
$C(t)$ = Concentration of Delayed Neutron Precursors
Λ = Prompt Neutron Generation Time
l_p = Prompt Neutron Lifetime
μ = Mean
σ = Standard Deviation
N = Rolling Period Sample Size

ACKNOWLEDGEMENTS

This work is supported by DOE, under DOE Idaho Operations Office Contract DE-AC07-05ID14517. Accordingly, the U.S. Government retains a nonexclusive, royalty-free license to publish or reproduce the published form of this contribution, or allow others to do so, for U.S. Government purposes. This research made use of Idaho National Laboratory's High Performance Computing systems located at the Collaborative Computing Center and supported by the Office of Nuclear Energy of the U.S. Department of Energy and the Nuclear Science User Facilities under Contract No. DE-AC07-05ID14517.

REFERENCES

1. E. Ryan, I. Trivedi, S. Terlizzi, A. Shields, M. Schanfein, G. Reyes, and R. Stewart, "Challenges and technology-driven opportunities for safeguarding microreactors," Annals of Nuclear Energy (2024).
2. S. Casey-Maslen, *The Treaty on the Non-proliferation of Nuclear Weapons*, Cambridge University Press, p. 67–97 (2021).
3. D. T. Consortium, "Definition of a digital twin," 2024. [Online]. Available: https://www.digitaltwinconsortium.org/initiatives/the-definition- of-a-digital-twin/.
4. E. Treviño, A. Shields, R. Stewart, J. Darrington, J. Scott, C. Pope, and C. Ritter, "Autonomous anomaly detection of proliferation in the AGN-201 nuclear reactor digital twin," Annals of Nuclear Energy, vol. 211, p. 110990 (2025).
5. R. Stewart, E. Treviño, A. Shields, K. Heaps, J. Darrington, Q. Williams, C. Pope, J. Scott, B. Baker, J. Palmer, B. Vainqueur, T. S. Palmer, C. Palmer, S. Bays, M. Schanfein, G. Reyes, and C. Ritter, "The AGN-201 Digital Twin: A Test Bed for Remotely Monitoring Nuclear Reactors," Annals of Nuclear Energy, vol. 213, p. 111041 (2025).
6. E. Schiassi, M. De Florio, B. D. Ganapol, P. Picca, and R. Furfaro, "Physics-informed neural networks for the point kinetics equations for nuclear reactor dynamics," Annals of Nuclear Energy, vol. 167, p. 108833 (2022).
7. C. Pope and W. Pheonix, Idaho State University AGN-201 Low Power Teaching Reactor - An Overlooked Gem, ser. Nuclear Reactors Spacecraft Propulsion, Research Reactors, and Reactor Analysis Topics (2021).
8. "Safety analysis report of the Idaho State University AGN-201m Research Reactor," Tech. Rep. License No. R-110, Docket No. 20-284 (2003).
9. M. Gorham, "Experimental parameterization of the Idaho State University AGN-201 Research and Training Reactor," master's thesis, Idaho State University (2012).
10. J. Leppänen, M. Pusa, T. Viitanen, V. Valtavirta, and T. Kaltiaisenaho, "The Serpent Monte Carlo code: Status, development and applications in 2013," Annals of Nuclear Energy, vol. 82, pp. 142–150 (2015).
11. R. Morris and J. Beaver, "openpointkinetics," (2018). [Online]. Available: https://github.com/richmorrison/openpointkinetics
12. X. Zhao, P. Li, K. Xiao, X. Meng, L. Han, and C. Yu, "Sensor drift compensation based on the improved LSTM and SVM multi-class ensemble learning models," Sensors, vol. 19, no. 18 (2019).
13. S. Hochreiter, "Long short-term memory," Neural Computation MIT- Press (1997).
14. G. McKenzie, T. Grove, and J. Hutchinson, "Using Prompt Neutron Decay Constant Measurements to Obtain Additional Kinetics Information," EPJ Web of Conferences 247, 09024 (2021).