INL/EXT-21-65074

Verification and Validation of Developed Short-Term Forecasting Models

November 2021

Cody Walker, Vivek Agarwal, Nancy Lybeck
Idaho National Laboratory

Pradeep Ramuhalli
Oak Ridge National Laboratory

Mike Taylor
Electric Power Research Institute

Idaho National Laboratory

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

DISCLAIMER

This information was prepared as an account of work sponsored by an agency of the U.S. Government. Neither the U.S. Government nor any agency thereof, nor any of their employees, makes any warranty, expressed or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness, of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. References herein to any specific commercial product, process, or service by trade name, trade mark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the U.S. Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the U.S. Government or any agency thereof.

INL/EXT-21-65074

Verification and Validation of Developed Short-Term Forecasting Models

Cody Walker, Vivek Agarwal, and Nancy Lybeck
Idaho National Laboratory
Pradeep Ramuhalli
Oak Ridge National Laboratory
Mike Taylor
Electric Power Research Institute

November 2021

Idaho National Laboratory
Idaho Falls, Idaho 83415

http://www.inl.gov

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

Page intentionally left blank

ABSTRACT

Recent advancements in machine learning (ML) and artificial intelligence (AI) technologies provide an opportunity for leveraging data-driven algorithms to predict future nuclear power plant (NPP) operating conditions by using recorded plant process data. Successfully implementing these models can lead to cost-reducing, conditioned-based predictive maintenance through optimized maintenance schedules and a reduction of unnecessary maintenance activities. This report discusses the verification and validation of short-term forecasting processes (i.e., data cleaning, feature selection, model optimization, and forecasting) developed in previous reports. The verification and validation (V&V) process demonstrates the expected precision and accuracy when the ML model encounters new datasets from different systems. Shapley additive explanations were used as the primary means of feature selection across these different data set. Individual models were trained for each data set, then validated through a cross-validation procedure. In this report, two different ML models were tasked to predict variables from three different plant process data sets with varying prediction horizons. The results indicate that support vector regression (SVR) outperformed long short-term memory (LSTM) neural networks in regard to each data set and each prediction horizon in this study, but further tuning and optimization could improve long short-term memory results. However, each forecasting model showed reduced performance as the prediction horizon was extended from 1 hour to 1 day ahead. Research is ongoing to evaluate the optimal input variable space, which is based on a given set of process parameters, to further improve forecasting accuracy.

Page intentionally left blank

CONTENTS

FIGURES

TABLES

Page intentionally left blank

ACRONYMS

AI	artificial intelligence
BWR	boiling-water reactor
DirRec	direct-recursive hybrid
FWCS	feedwater and condensate system
FDWP	feedwater pump
FWPT	feedwater pump turbine
LSTM	long short-term memory
ML	machine learning
NPP	nuclear power plants
PWR	pressurized water reactor
RMSE	root mean square error
SHAP	Shapley additive explanations
SVR	support vector regression
V&V	validation and verification

Page intentionally left blank

1 INTRODUCTION

Nuclear power is the single largest contributor of non-greenhouse-gas-emitting electric power generation in the United States, therefore it is imperative to ensure its continued operation [1]. But despite excellent safety records, premature closures in the nuclear industry have occurred due to economic concerns. The largest cost to legacy light-water plants are the operating and maintenance (O&M) costs. The overall goal of the project is to develop and demonstrate predictive maintenance using online monitoring of critical systems. By predicting a component's condition in the future, maintenance can be scheduled and performed in a timely manner, leading to a reduction in maintenance costs associated with unnecessary maintenance, parts and labor, and unplanned, forced, or extended outages.

Nuclear power plants (NPP) record large volumes of heterogenous data from various equipment and systems. This data typically includes plant process parameters, maintenance records, technical logs, online monitoring data, and equipment failure data. NPP data in conjunction with machine learning (ML) and artificial intelligence (AI) technologies can leverage this data to provide actionable information on future conditions of critical equipment and systems. The feature selection and forecasting methodologies of these techniques should be generalizable to ensure the most accurate and consistent calculations when applied to data sets. The developed models and methodologies should then be validated using an independent data set to build trust in the applicability of these models and methodologies in new environments.

This report is the latest in a series of documents detailing online monitoring, wireless communication networks, and the development of diagnostic and prognostic models using heterogenous data regarding critical balance of plant equipment in NPPs [2-7]. The previous documents in this series evaluated potential wireless technologies and communication networks inside NPPs, based on their performance and economics [2-5]; assessed vibration sensors with wireless capabilities [6]; developed formal methodologies for cleaning data and objectively comparing ML prognostic models [7]; and development of condition-based monitoring and preventative-maintenance optimization [8].

In previous reports [7,8], the focus was aimed at the condensate and feedwater system for a boiling water reactor (BWR). This system contains many interconnected components including condensate pumps, condensate booster pumps, feedwater heaters, condensers, and more. Three ML models (long short-term memory networks, support vectors, and random forest) were used to predict future plant parameters under steady-state operating conditions. With the components being interconnected, their parameters can interact or be influenced by other components, so feature selection becomes important in selecting the pertinent plant parameters for making a prediction.

The goal of this report is to demonstrate the expected accuracy and consistency of the developed short-term forecasting ML methodologies when given new data from different systems. Verification and validation (V&V) is an important exercise to demonstrate the strength and applicability of research results. Application of the developed techniques and methodologies from previous research to new, independent data sets allows for a better understanding on how the models might respond in a new environment or when formally deployed. This report expands the focus to include both a pressurized water reactor (PWR) steam generator system and PWR turbine system. One model was not created to predict each system, but rather, the V&V process focused on how well the methodologies (i.e., data preprocessing, feature selection, model optimization, and forecasting) could be applied to new data sets. This report focuses on the V&V of the short-term forecasting models that are tasked with predicting the component's future condition that could indicate component's health given the operating conditions have not changed, demonstrating the process by evaluating how accurately for two ML models, long short-term memory neural networks (LSTM) and support vector regression (SVR), perform for three independent data sets with two different prediction horizons (1 hour and 1 day).

Section 2 discusses the systems of interest and the methodologies used for short-term forecasting. Section 3 discusses the V&V results – including the differences in the independent data sets, the data cleaning and preprocessing, and the performance of the short-term forecasting models. Section 4 concludes with a brief summary and details a path forward for future research.

2 SHORT-TERM FORECASTING MODELS

NPPs are interconnected systems composed of many subsystems (e.g., reactor, steam generators, turbines, etc.) and each of these subsystems are composed of numerous components (e.g., motors, pumps, pipes, valves, and bearings). Each of these systems and many of their components are being monitored either through online monitoring or are manually inspected periodically. The type of data collected is application dependent and may include temperatures, pressures, flows, inches of expansion, gross load, reactor thermal load and differential pressures. This wide variety of components, sampling frequencies, data types, and locations throughout the plant leads to a heterogenous source of data that must be handled appropriately before use both in general and specifically for ML modeling. Data from each subsystem or set of components may be different and identifying a single set of features for comparing prediction accuracy is impractical. However, the procedures for data cleaning, feature selection, and forecasting can be applied to each data set and the overall prediction accuracies can be compared. Then, an expected prediction accuracy can then be developed based on the data acquired, feature selected, and forecasting methodology applied. This section covers the three systems of interest used as the data sources in this study, the feature selection process, and the short-term forecasting models.

2.1 Systems of Interest

For the V&V process, the same methodologies and workflow were performed on three independent data sets. The first of which came from a boiling water reactor (BWR), while the other two came from pressurized water reactors (PWR). The specific systems of interest include a boiling water reactor's condensate and condensate booster pump, a pressurized water reactor's steam generator feedpump, and another pressurized water reactor's main turbine. The fundamentals of each component and the specific signals recorded are explained below.

2.1.1 BWR Condensate and Condensate Booster Pumps

Condensate and condensate booster pumps are found with the condensate and feedwater system. The primary purpose of this system is to condense steam and collect the drainage in the main condenser before purifying, preheating, and pumping the water back to the reactor vessel [9]. The condensate pumps provide the driving force for pushing the condensate through auxiliary systems such as the steam jet air ejectors condenser, steam packing exhaust condenser, off-gas condenser, and demineralizers - all of which work to condition the condensate. Afterwards, the condensate booster pumps are the driving force of the flow as the condensate travels through a string of low-pressure heaters that work to preheat the water to the correct temperature. In the BWR system of interest, the condensate and condensate booster are driven by a shared motor.

2.1.2 PWR Steam Generator Feedpump

In a PWR system, the primary coolant flows from the reactor to the steam generator to transfer its thermal energy to the secondary coolant through many tubes [10]. With sufficient heat, the secondary coolant starts to boil into steam before being sent to the turbine system. The steam generator feedpump is what supplies the driving force for the secondary side's coolant to flow through the steam generator.

2.1.3 PWR Turbine

In a PWR system, the steam formed in the steam generator is passed to the main turbine generator where it is used to make electricity [10]. The steam is then directed to low pressure turbines before being routed to the main condenser. Throughout this process, the steam is piped through moisture separators and reheaters (MSRs). The MSRs dry and reheat the steam to prevent potential turbine damage due to the moisture content of the steam while also improving the efficiency of the turbine.

2.2 Feature Selection

The same data cleaning and feature selection techniques were used for each dataset to create a fair basis for comparison. The data cleaning process is discussed in detail in Section 3.2 along with the discussion of the recorded process parameters. The primary technique used for feature selection was the Shapley additive explanations (SHAP) approach.

SHAP has been used to improve model interpretability, but it can also identify important variables by how they contribute to the ML model output [11,12]. SHAP is classified as an additive feature attribution method, which means the output of the model is a linear addition of the inputs as seen in Equation 1,

where f(x) is the original model, g(x') is the explanation model, ϕ_0 represents the model's baseline estimate, M is the number of input variables, is the feature attribution values, and are the input variables themselves [13]. The baseline, ϕ_0, is determined by averaging over all predictions for a given model [14]. The SHAP approach uses a method akin to game theory to estimate the feature attribution values. Input variables are grouped, and the models are trained. By doing this in several different combinations, the influence of each variable over the model's output can be determined. Input features that contribute more to the model's output have larger feature attribution values associated with them. In this way, the feature attribution values are related to the feature's importance as larger attribution values (either positive or negative) will lead to larger changes in the model's output. The feature attribution values can be approximated through a variety of methods depending on the model being use, such as Kernel SHAP for support vector machines, Deep SHAP for deep learning techniques, and Tree SHAP for decision trees [11]. Feature selection is made by retaining only variables with larger attribution values and discarding those associated with smaller attribution values. One of the drawbacks with the SHAP method is the computational load that expands depending on both the type of explanation model (e.g., Tree, Kernel, or Deep SHAP) and the number of variables being analyzed. It could be useful to reduce the number of variables with correlation analysis before implementing the SHAP method to reduce redundant variables.

2.3 Short-term Forecasting Models Selected

Two models were selected for short-term forecasting based on previous performances [6]. These models were the Long Short-Term Memory (LSTM) neural network and Support Vector Regression (SVR).

2.3.1 Long Short-Term Memory Networks

LSTMs are a recurrent-type neural networks that were specifically designed to learn and encode the long-term relationships between the inputs and outputs through via unique memory cells that. These memory cells store, within the network, a hidden state within the network that interacts with the current input. The memory, or stored state, is updated as new inputs become available. A "forget" gate within the network determines which information to forget from the previous memory cell during the next iteration as a hidden output, which improves performance. LSTMs excel at utilizing temporal information to determine long-term relationships between the input and outputs. For more illustration on the construction and implementation of LSTMs, a guide has been provided by Greff [15].

2.3.2 Support Vector Regression

SVR is a kernel-based regression technique that makes estimations based on a selected subset of training data called support vectors. Support vectors are chosen based on a kernel function and an optimization routine. The kernel function first transforms the data into a higher-dimensional space so that the relationship between variables is assumed to be linear. Commonly used kernel functions include radial basis functions, polynomial, and sigmoidal functions. The goal of the optimization routine is to then find a hyperplane that minimizes generalization errors, which is the sum of the training error and the confidence level. SVR is unlike other regression techniques as training data within a selected tolerance (epsilon-tube) does not penalize the loss function during optimization. The training data that lies closest to the hyperplane is then selected as the support vectors. A more detailed guide to SVR was provided by Smola [16].

2.4 Multi-step Forecasting

The simplest forecast is one-step ahead since it represents the least amount of extrapolation from the known data. Intuitively, we expect better results as there are fewer opportunities for unseen transients, depending on the size of the step. The 24-hour ahead predictions in the previous report [7] were created using one-step ahead predictions. This was accomplished by resampling the data so that each time step was 24 hours apart rather than the original 1-hour frequency. This method yielded acceptable results for that dataset, primarily due to the amount of data available – 5 years' worth. However, the other two V&V datasets discussed in this report only contain a year's worth of data. By resampling the data from an hourly to daily, and then further dividing the data into training, testing, and validation sets, the V&V dataset becomes sparse for learning trends. This resampling strategy is therefore inappropriate for these datasets. With limited data, each point is critical for training the model. Several methods for making multi-step ahead time-series predictions were tested to see the optimal strategy for forecasting. These prediction strategies include direct multi-step forecasting, recursive multi-step forecasting, and direct-recursive hybrid multi-step forecasting [18]. These methods are discussed in further detail in the subsections below.

2.4.1 Direct Multi-step Forecasting

The first method for making predictions that are multiple time steps ahead is direct multi-step forecasting method. In this method, the model is trained to directly predict a set time into the future. This method is similar in appearance to a one-step ahead prediction, but with an increased prediction horizon. Equation 2 shows an example of a one-step ahead model,

where a model uses the previously observed data, x_t, x_{t-1}, …, x_{t-n}, to make a prediction that is a single time step ahead, x_{t+1}, and n is the number of previously observed time steps to be inputted into the model. In direct multi-step forecasting, this prediction horizon of one-time step is extended to the desired length. Equation 3 shows an example of a model predicting two-steps ahead with the same previously observed data as Equation 2. The difference between equation 2 and 3 is the model's predicted outcome (x_{t+1} compared to x_{t+2}) due to the extended prediction horizon.

This prediction horizon can be extended further. In this report, hourly recorded data were used to predict component measurements twenty-four time steps, or one day ahead. However, there are limitations to this technique. Extending the prediction horizon will, in most cases, diminish the accuracy of the predictions. For this type of model, there is no opportunity for the model to learn the dependencies between forecasted predictions. As one variable changes, it may directly affect another, since this type of model might overlook this type of interaction as it directly predicts future values.

Similarly, a direct, multi-output, multi-step forecast casting strategy can be implemented in which one model attempts to predict each consecutive future step, as seen in Equation 4.

This type of model will have similar or exacerbated drawbacks as direct multi-step forecasting due to the additional predicted outputs. This will likely lead to a more complex model that will require more time to train as well as more data to prevent overfitting. So as an overview: Model 1, in equation 2, attempts to predict one-step ahead; model 2, in equation 3, attempts to predict 2 steps ahead using the same inputs as model 1; and model 3 tries to predict both the one- and two-steps ahead while still using the same input parameters.

2.4.2 **Recursive Multi-step Forecasting**

Another method for making multi-step ahead predictions is recursive multi-step forecasting, which entails a single model that makes one-step ahead predictions. After a prediction is made, the input window is shifted by one step, and the model's own prediction is used as one of the inputs. This process of recursively using the model's output as an input is repeated until the desired prediction horizon is reached. Equation 5 and 6 shows a single step in the recursive multi-step forecasting process. In Equation 5, model1 makes a single one-step ahead prediction, $x_{(t+1)}$. This prediction is then fed back into itself, as seen in Equation 6, as model1 is then used to make another one-step ahead prediction, $x_{(t+2)}$.

This model has some disadvantages, namely, the accumulation of prediction errors. There will inevitably be some prediction error between the predicted output of Equation 5 and the ground truth. This error is then inserted into the model again in Equation 6, where the error can increase the deviations from the ground truth, thus accumulating predicting errors. Longer prediction horizons lead to more accumulation of errors.

2.4.3 **Direct-Recursive Hybrid Multi-Step Forecasting**

Direct-recursive hybrid (DirRec) multi-step forecasting is a combination of direct and recursive forms of prediction to mitigate the disadvantages from each method. In DirRec multi-step forecasting, several models are trained, each with a minimal prediction horizon. However, the prediction from one model is then fed as an input to another model that makes the same length prediction. In this way, a series of models are used in succession to extend the prediction horizon. The first model, seen in Equation 7, is used to predict a single step ahead, $x_{(t+1)}$. This prediction is then used as an input for model 4, seen in Equation 8, which is used to make another single step ahead prediction, $x_{(t+2)}$.

This approach aims to minimize the disadvantages from both the direct and recursive approaches. With a series of models, the dependencies between predicted outcomes can be observed, thus reducing the disadvantage of direct approaches. With multiple models being created, there is also an opportunity to correct some of the accumulating errors that one specific model could make, thus reducing the recursive model drawbacks. However, creating numerous models can be computationally expensive to train. With each model making one-step (one hour) ahead predictions, twenty-four models would need to be trained to estimate one day in advance.

In this report, the recursive multi-step forecasting method had prohibitively large accumulations of prediction errors. And although the DirRec multi-step forecasting method avoided such error accumulation, the training of multiple models became computationally expensive. A sample case from each forecasting method is shown in the Section 3.3 to demonstrate its forecasting ability. The direct multi-step forecasting method was chosen as the primary method for forecasting in this report due to its relatively low prediction error and computational expense.

2.4.4 **Cross Validation and Model Comparison**

Comparisons across independent data sets can indicate how models or methods will respond in new environments or when they are deployed formally. With the systems of interests and features of the data sets being so diverse in this report what is primarily being compared is the methodology for generating predictions based on the raw data. LSTM and SVR models were trained for each dataset then used to forecast both one-step and twenty-four steps into the future. A ten-fold cross validation technique was used to resample the data into ten different partitions. One partition was withheld at a time, while the remaining nine were used for training and validation. The withheld partition was then used for testing to evaluate the model's performance on an unseen set of data. The models would then be evaluated on its average performance over all iterations. With the data processing, feature selection, and model creation being the same for each data set, the cross-validation results should provide insight into how this methodology might respond to new environments.

3 VERIFICATION AND VALIDATION RESULTS

This section provides an overview of the results achieved when applying the specific methodology described above to the operational data from three different NPP subsystems. The primary objective is to analyze and compare the capabilities and limitations of the data-driven, short-term forecasting models as they predict future plant data one- or multiple-time steps into the future. Each data set undergoes the same set of data processing and feature selection before predictions are made using LSTM and SVR models. The relative accuracy of each model, as determined using ten-fold cross validation, is compared for each dataset.

3.1 Data Description

The available process data comes from three different subsystems within three separate NPPs. These systems include a feedwater and condensate system (FWCS) of a BWR, a steam generator feedpump from a PWR, and a PWR main turbine and each varied in the amount of data collected on them, the sensors used, and the overall operating conditions.

3.1.1 BWR Condensate and Condensate Booster Pumps

The available process data recorded for FWCS system in the BWR include:

- generator gross load,
- average feedwater flow,
- feedwater pump header and discharge pressures,
- feedwater temperatures,
- condensate and booster pump header and discharge pressures,
- condensate pump motor currents,
- condenser hotwell level,
- turbine exhaust temperatures,
- condensate and booster pump temperatures,
- drive motor temperatures.

This data from the plant historian covers five years' worth of operation and was downsampled to an hourly frequency before being provided by the plant. During the five years of operation, there were multiple planned refueling outages, periods of steady-state operation, reactor trips, and derates of varying severity. Some maintenance work order information was available, but not enough to determine the cause of each derate. Discussions with the plant engineers indicated that a majority of the derates were cause by systems other than the FWCS (i.e., no FWCS failures were seen in any components within the system during the recorded time period). There was no access to the signals outside of the FWCS system to investigate those derates further. For this data set, the prediction variable is a bearing temperature inside a condensate pump within the FWCS system. Bearing temperatures can be indicative of the pump's condition.

It is also worth noting that the average plant gross load from this data set differed by about 0.5% - 1% from one refueling cycle to another. Models trained solely on cycles with similar power outputs produced less accurate results when tested on cycles with a different power level as those conditions were not within the training scope. By ensuring sufficient diversity in the training data set, this source of bias was eliminated.

3.1.2 PWR Steam Generator Feedpump

The available process data for the PWR steam generator system includes:

- the suction and discharge pressures,
- the feedwater pump speed,
- feedwater pump bearing temperatures,
- feedwater pump oil drain temperatures,
- feedwater heater temperatures,
- feedwater flows,
- the thermal reactor power.

This data was also provided at an hourly frequency. The data covered one operating year during which there was a refueling outage that lasted for roughly one month. Information on maintenance work orders is available for this time period, but no degradation related to this subsystem was recorded in the maintenance work order files. For this dataset, the flow moving through the steam generator feedpump was predicted. This flow can be used to determine if there are leaks present in the loop.

3.1.3 PWR Turbine

The available process data related to another PWR's main turbine included:

- the generator gross load,
- pump speeds,
- bearing temperatures,
- oil drain temperatures,
- feedwater flows,
- bearing vibrations,
- turbine eccentricity,
- turbine differential expansions,
- ambient air temperatures.

The available data covered one year's of steady-state operation with some derates, but no refueling outages. The information was collected at an hourly frequency, but no associated maintenance work order information to determine ground truths relating to component condition or cause of derates. For this PWR, both the main turbine front bearing temperature and the generator's output were predicted. The turbine's condition and the generator output can both be linked to the plant's performance.

3.2 Data Cleaning and Processing

The time series data for each data set was cleaned, and the features were selected before the data was separated into steady-state and non-steady-state condition categories. The features in this report are the plant parameters themselves (i.e. temperature, pressure, etc.). Steady-state conditions consisted of all plant operations when the reactor's thermal output was above 90% of its nominal value. Variability could be seen in the reactor's thermal output during steady-state operation, but major derates, ramp downs, trips, and refueling outages were excluded from this analysis.

The first step in cleaning the data was to address all the missing values, most of which corresponded to when the component was offline. These values remained as is. If the component was online, then the missing value was most likely recorded due to an error in the sensor or during archival. These missing values were replaced via interpolation using its nearest neighbors. Other missing values were seen with daylight savings time as periods of time were “skipped” once per year. These times were filled in as with the previous values.

The second step is to address any outliers found within the data. Outliers were selected based on being four standard deviations away from the average. Many of the signals are heavily skewed in one direction during steady-state operation, so three standard deviation flagged as much as 1% of the data, which seemed high. Four standard deviations flagged the number of outliers to be about 0.3% of the data. These outliers were then replaced with a median filter of size 51, which corresponds to just over 2 days’ worth of data. It is assumed that these outliers were not due to operational changes in the plant.

After the data cleaning, the variables were scaled to a standard normal distribution with a zero mean and unit standard deviation. Features for each model were then selected based off their feature importance using SHAP values. The feature importances were calculated as the mean SHAP values for their respective variables. A set threshold for feature importance was not selected, but rather a dynamic one. Features within an order of magnitude of the largest feature importance value were selected as input variables. For example, in Table 1, when selecting input variables to predict bearing vibration, excluding the vibration signal itself, the largest feature importance is the turbine speed with a value of 0.608. This feature being the most important is logical because a higher speed can subsequently lead to higher levels of vibrations as the system is worked harder. The final value shown in Table 1 is the turbine shell expansion with a value of 0.029. This value is more than one order of magnitude away from the highest value, so the shell expansion and all variables with lower feature importance values were excluded as model inputs.

Table 1: Feature selection for predicting bearing vibration based on SHAP importance values.

Feature	Importance
Turbine Speed	0.608
Turbine Rotor Expansion	0.309
Turbine Differential Expansion	0.161
Flow to Turbine	0.151
Shell Expansion	0.029

The variables included in this analysis as they had the largest SHAP values are shown in Table 2. Many variables are relatively slow moving in time (e.g., temperature), so it is advantageous for the model to know the variable’s value at previous time steps. Input variables used to forecast the main turbine thrust bearing temperature include the temperatures of other bearing that were in close physical proximity to itself. To predict generator output, the other input variables included flow to the feedwater pump’s turbine (FWPT), the feedwater pump (FDWP) bearing temperature, and the turbine’s shell expansion. Flow to the FWPT had a negative correlation to generator output. As more flow was redirected to the FWPT, less output from the generator was seen. However, generator output had a positive correlation with both the FDWP bearing temperature and the turbine shell expansion. As the generator output increased, similar increases were seen in the bearing temperature and shell expansion as the system ran harder and hotter. To predict steam generator feed flow, the other input variables included steam generator (SG) flow from the other loop, the steam generator feed pump (SGFP) bearing temperature, and the feedwater (FW) heater outlet temperature. The condensate pump temperature just required itself and another condensate pump bearing temperature as input variables.

Table 2: List of predicted variables and the variables used to predict them as selected based on their SHAP values.

Variable Predicted	Other Input Variables		
Main Turbine Thrust Bearing Temperature	Front Thrust Bearing Temp	Rear Thrust Bearing Temp	
Generator Output	Flow to FWPT	FDWP Bearing Temp	Turbine Shell Expansion
Steam Generator Feed Flow	SG flow loop 2	SGFP Bearing Temp	FW Heater outlet Temp
Condensate Pump Temperature	Condensate Pump 2 Bearing Temp		

3.3 **Short-term Forecasting Model Performance**

This report focuses on how well LSTM and SVR predicted plant process data sets over two prediction horizons (1 hour and 1 day). Three different multi-step approaches were tested for the main turbine thrust bearing temperature to determine which approach would be used for all other predictions. The direct, recursive, and DirRec multi-step approaches were used to estimate the bearing's temperature 24-steps into the future using the LSTM model, and the results were compared using the root mean square error (RMSE) in the prediction as calculated in Equation 9,

where N is the total number of predictions, is the model's predicted output, and is the observed output.

The direct, recursive, and DirRec multi-step approaches produced RMSE's of 0.294, 4.816, and 0.301 for the anonymized bearing temperature dataset, respectively. For this model, the recursive method results that that were considered poor due to the much larger RMSE than seen for the other two methods. This larger RMSE was due to the accumulation of prediction error while recursively looping the output from the one-step ahead model back into the input. The predictions using the direct and direct-recursive methods on the anonymized temperature data can be seen in Figure 1. The direct method, seen in green, more closely follows the peaks as the temperature of the bearing oscillates. The cause of this oscillation, which is on the order of 1-2 degrees Celsius, is unknown due to a lack of information on how the unit was operated or if the system had any underlying conditions. Although the DirRec method had comparable results, the direct method was chosen based on having the lowest RMSE and was used for the remaining of the V&V process.

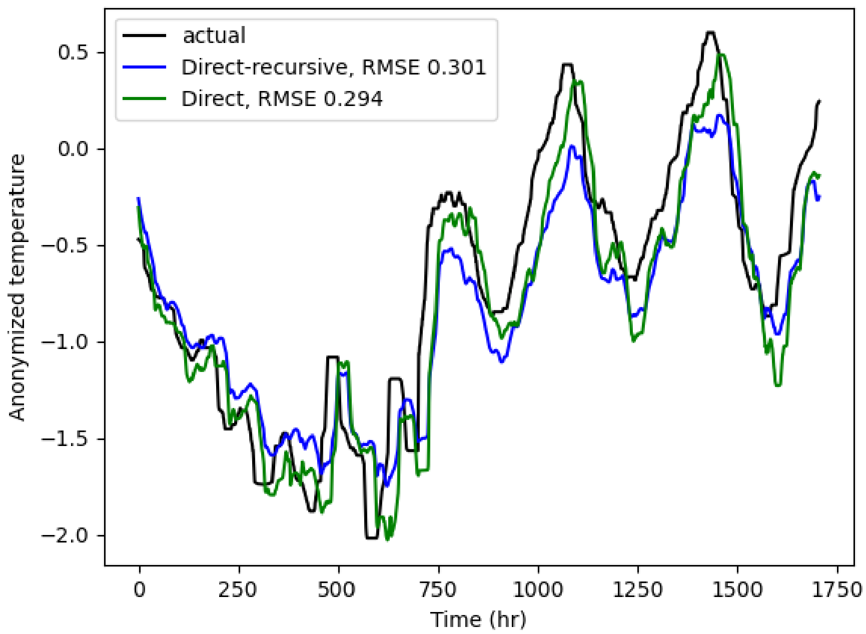

Figure 1: Comparison on direct versus direct-recursive forecasting strategies.

The hyperparameters for LSTM are important components that help determine the quality of the model's predictions. Optimal LSTM hyperparameters were identified using a grid search method using the BWR's condensate pump bearing temperature data set, and the best results for the hyperparameters were used for all subsequent LSTMs. The chosen hyperparameters can be found in Table 3. The model was trained for a maximum of 250 epochs or until the validation loss did not improve for 6 epochs. Early stoppage of training was used to prevent overfitting the model to the training data. The dropout layer and L1 & L2 regularizers were all added to improve the LSTM's robustness by reducing overfitting. Only a single LSTM layer was used in this model.

Table 3: Optimal hyperparameters used for the LSTM models.

Hyperparameter	Value
Number of units	1000
Batch size	64
Epochs	up to 250
Dropout	20%

Validation split	10%
Optimizer	Adam
L1 & L2 regularizer	1.0E-05

LSTM and SVR were used make one-hour (i.e. one step) and one-day (i.e. 24 step) ahead predictions Data from three different nuclear plants (i.e., PWR 1, PWR 2, and BWR). PWR 1's and PWR 2's data sets contained one-year of data, while the BWR's data set contained five-years of data. The mean of the RMSEs from the cross-validation approach demonstrates the expected accuracy of the modelling method when used to make predictions. The RMSE's means and standard deviations from the LSTM and SVR models for the one-hour ahead and one-day predictions are shown in Table 4.

Table 4: 1-hour and 1-day ahead predictions for both LSTM and SVR.

Data set			1-step ahead		24-steps ahead	
Plant	Parameter Predicted	Model	Mean RMSE	Std Error	Mean RMSE	Std Error
PWR 1	Main Turbine Bearing Temp	LSTM	0.0796	0.0411	0.7932	0.5450
		SVR	0.0214	0.0080	0.3194	0.1202
PWR 1	Generator Output	LSTM	0.2871	0.2031	2.2636	2.8338
		SVR	0.0806	0.0422	1.5611	1.2424
PWR 2	Steam generator flow	LSTM	2.4792	3.0455	12.435	17.333
		SVR	1.4070	2.4270	5.6299	5.3154
BWR	Condensate Pump Bearing Temp	LSTM	0.0792	0.0722	0.2991	0.2724
		SVR	0.0323	0.0496	0.2238	0.2184

In Table 4, the RMSEs were calculated after converting the predictions back to their respective engineering unit of measurement. Temperatures were captured in Celsius, generator output in megawatts, and steam generator flow in thousands of gallons per minute. The first point of comparison comes between the prediction accuracy of the SVR compared to the LSTM. In each dataset, the SVR outperformed the LSTM. This may be due to how each model was optimized. SVR has a convex optimization, so choosing the optimal hyperparameters (i.e., regularization parameter, epsilon-tube, and kernel coefficient) can be done efficiently. For LSTM, there are many hyperparameters to consider, some of which are listed in Table 3, which makes the optimization more difficult and computationally expensive. After the hyperparameters are chosen, the LSTM model still must be trained to the data set of interest. In this report, the optimal hyperparameters were found for one data set then applied to each other data set. The BWR condensate pump bearing temperature data set, in which the LSTM hyperparameters were optimized, had the closest results, in terms of mean RMSE, to the SVR out of all the data sets. This suggests that further improvements may be seen if the optimal architecture was selected for each LSTM with each new data set.

Comparing results within Table 4, there is a trend that predicting 24-steps ahead performs worse than predicting 1-step ahead. This logical as more events and transients can occur over a longer prediction horizon, which makes the forecasting more difficult. Another trend that can sometimes be seen is a lag in the 24-step forecasts, as seen in Figure 2. The 1-step ahead predictions for both the LSTM and SVR seem to trend well with the actual measurements; however, the 24-step ahead predictions produce a noticeable delay. The predictions mimic the same patterns as the temperature fluctuations but does not accurately predict them as they occur. This is most likely due to a lack of predictor information that would inform the model of the cause of these one-to-two-degree fluctuations that appeared to oscillate over a two-week period. The source of the fluctuations was not determined due to a lack information such as operating or maintenance logs. With derates being removed, the 24-step ahead model was still expected to predict 24-steps ahead even though this no longer corresponded to the original 24-hour intent. This could be one cause of the inflated mean RMSE values.

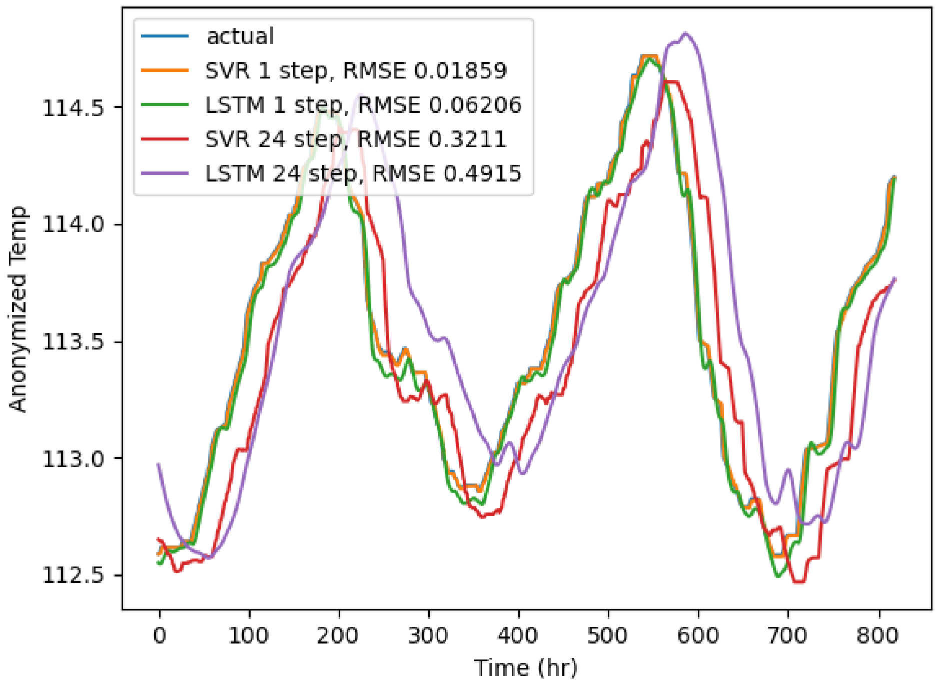

Figure 2: Forecasting the main turbine thrust bearing temperature using LSTM and SVR for both one- and 24-steps-ahead predictions.

Another interesting comparison can be seen within the LSTM's prediction of PWR 1's main turbine bearing temperature from 1-step to 24-step when compared to the LSTM's prediction of BWR's condensate pump bearing temperature from 1-step to 24-step. The mean RMSEs for both 1-step ahead predictions are very close to one another with 0.0796 and 0.0792, respectively. However, the 24-step ahead predictions deviate significantly with mean RMSEs of 0.7932 and 0.2991, respectively. The difference seen between the 24-hour prediction mean RMSEs can most likely be attributed to the increased amount of data that the BWR data set had (5 years versus 1 year) that yielded an increase in the prediction accuracy. Although the LSTM still did not beat the SVR in the overall accuracy, the LSTM did see a bigger improvement when moving from 1-step to 24-step ahead predictions for this variable.

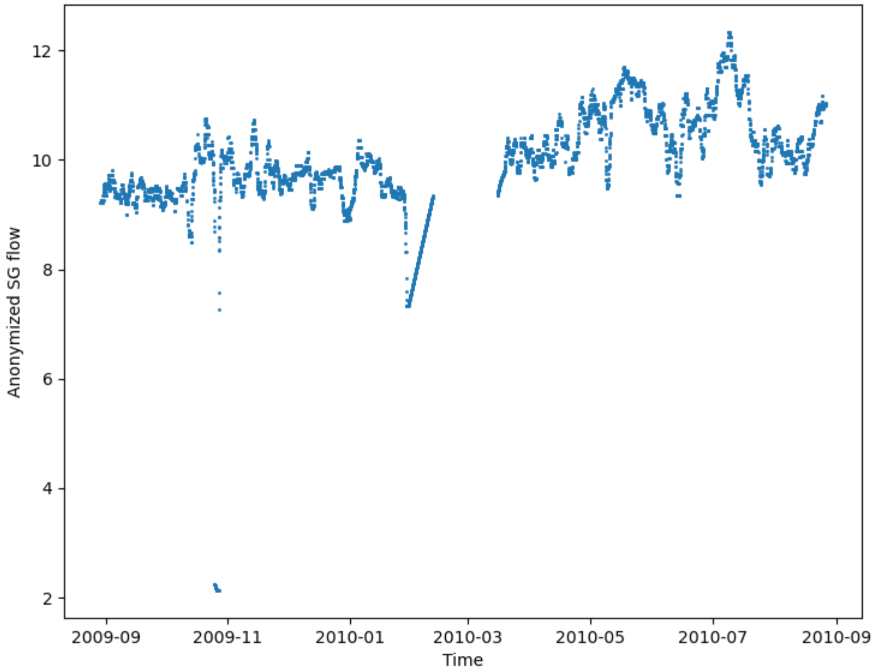

Figure 3: Anonymized steam generator flow and time that shows a significant derate at 2009-11 and an outage at 2010-03.

For both the models and prediction horizons in Table 4, the RMSE standard deviation for PWR 2's flow predictions are large. This is due to the fact that two separate sections of the flow data, which is experiencing abnormal behavior, shown in Figure 3 where the y-axis was anonymized, and the x-axis was shifted to protect the plant's identity. In Figure 3, the first large deviation can be seen at time 2009-11. This was a large reduction in the flow that occurred when the thermal power of the reactor was above 90%. The other section of abnormal behavior occurred around 2010-02, where the reactor was being to ramp down toward a regularly scheduled outage that lasted through 2010-03. The flow appears to linearly increase in a strict fashion. Each of these sections were poorly predicted as their behavior were not captured in the training data. With these two, operational anomalies present, the mean RMSE for the SVR 24-step ahead prediction was 5.629. After removing these outliers from the testing data, the mean RMSE was reduced to 1.744. In the data cleaning process, each of these abnormal sections were missed. The large derate in flow remained within 3 standard deviations, and the linearly increasing section occurred while the reactor thermal power was above 90%.

Overall, each model predicted performed well based on the new data sets. However, as seen with the two flow outliers, data-driven models can only effectively be relied on to predict activity that has been seen within the training data set. The model performance can be improved by supplementing the training set with more information if available, although unfortunately, this is rarely the case.

4 SUMMARY AND PATH FORWARD

This report demonstrated the expected accuracy and consistency of short-term forecasting ML models when given new data from different systems. Prognostic and short-term forecasting models are a key enabling technology that can be used to reduce O&M costs by enabling condition-based predictive maintenance. This research showed that a model's prediction accuracy depends on several factors, including the model's architecture, the prediction horizon, and the training data set itself. Each ML model needs to be optimized in some way to the data set of interest as the model's performance directly relies upon it. SVRs have fewer hyperparameters than LSTMs, which makes it less computationally expensive to optimize. This research also showed that model performance degraded as the prediction horizon increased. In addition. this research showed the importance of having a representative training set. When the plant experiences transients not within the data set, as seen in the PWR 2 flow predictions, the model cannot be relied on to predict these unseen transients accurately. The analysis in this report only covered operating conditions at steady-state conditions as defined by the reactor's thermal power. There is a need to expand this work to non-steady-state operating conditions in order to cover the full range of the reactor's operating life cycle (i.e., derates, trips, and ramp downs towards outages). During these transient operating conditions, the system is still vulnerable to degradation and the knowledge of critical operating equipment's future condition is essential for enabling predictive maintenance.

This report also described the SHAP feature importance values for selecting a representative data set based on the feature importance. There are other feature selection techniques (i.e., filter, wrapper, and embedded methods [17]) and dimensionality-reduction methods (e.g., principal component analysis) that should be explored to see if an optimal input variable space can be developed. Feature importance values can potentially determine the features that contribute the most when predicting future plant parameters. This could incredibly useful when predicting certain parameters, such as gross load. In the results section, LPE flow was one of the major influencers for forecasting gross load. As low-pressure steam is exported to other areas of the plant, the gross load is reduced. With a wide array of sensors across many components and systems, new critical components can be determined based on how they influence other system parameters. These components would also then need to be analyzed by subject matter experts to determine the extent of the causal relationship.

5 ACKNOWLEDGEMENTS

This report was made possible through funding by the U.S. Department of Energy's Office of Nuclear Energy under the Nuclear Energy Enabling Technologies Program. We are grateful to Melissa Bates and Daniel Nichols at the U.S. Department of Energy and Pattrick Calderoni at Idaho National Laboratory for championing this effort. The data used in this research was provided by Exelon Generation Company. We are also thankful to the plant engineers for their technical discussions on the data.

6 REFERENCES

1. Office of Nuclear Energy, "Light Water Reactor Sustainability (LWRS) Program," Accessed on November 1, 2021: https://www.energy.gov/ne/nuclear-reactor-technologies/light-water-reactor-sustainability-lwrs-program.
2. V. Agarwal. 2020. "Wireless Sensor Modalities at a Nuclear Power Site to Collect Vibration Data," Idaho National Laboratory, Idaho Falls, ID. INL/EXT-20-58548.
3. K. A. Manjunatha and V. Agarwal. 2019. "Techno-Economic Analysis Framework for Wireless Networks in Nuclear Power Plants," Idaho National Laboratory, Idaho Falls, ID. INL/EXT-19-55830.
4. V. Agarwal. 2021. "Identification of Balance of Plant Asset and Wireless Instrumentation to Enable Predictive Maintenance," Idaho National Laboratory, Idaho Falls, ID. INL/EXT-21-54866 Rev 1.
5. Kosuhik Manjunatha and Vivek Agarwal, "Multi-Band Heterogeneous Wireless Network Architecture for Industrial Automation: A Techno-Economic Analysis," Wireless Personal Communication, Accepted for Publication.
6. K. A. Manjunatha and V. Agarwal. 2019. "Review of Wireless Communication Technologies and Techno-Economic Analysis," Idaho National Laboratory, Idaho Falls, ID. INL/EXT-19-53966.
7. P. Ramuhalli, C. Walker, V. Agarwal, and N. Lybeck. 2020. "Development of Prognostic Models Using Plant Asset Data," Oak Ridge National Laboratory, Oak Ridge, TN.

8. C. M. Walker, V. Agarwal, N. J. Lybeck, P. Ramuhalli, and M. Taylor. 2021. "Nuclear Power Fault Diagnostics and Preventative Maintenance Optimization," Idaho National Laboratory, Idaho Falls, ID. INL/EXT-21-61427.
9. U. S. Nuclear Regulatory Commission Technical Training Center. 1998. "BWR/4 Technology Manual (R-104B)," Ml022830867, pp. 1-442. https://www.nrc.gov/docs/ML0228/ML022830867.pdf.
10. U. S. Nuclear Rregulatory Commission. 2007. "Reactor Concepts Manual: Pressurized Water Reactor (PWR) Systems," 603:28. https://www.nrc.gov/reading-rm/basic-ref/students/for-educators/04.pdf.
11. S. M. Lundberg and S. I. Lee. 2017. "A Unified Approach to Interpreting Model Predictions," Advances in Neural Information Processing Systems 30, December (Section 2), pp. 4766–4775.
12. S. Mangalathu, S. H. Hwang, and J. S. Jeon. 2020. "Failure mode and effects analysis of RC members based on machine-learning-based SHapley Additive exPlanations (SHAP) approach," *Engineering Structures*, 219:110927. doi: 10.1016/j.engstruct.2020.110927.
13. S. M. Lundberg, G. G. Erion, and S. I. Lee. 2018. "Consistent individualized feature attribution for tree ensembles," 31st Conference on Neural Information Processing Systems (NIPS 2017).
14. A. L. Booth, E. Abels, and P. McCaffrey. 2021. "Development of a prognostic model for mortality in COVID-19 infection using machine learning." *Modern Pathology*, 34(3):522–531. doi: 10.1038/s41379-020-00700-x.
15. K. Greff, R. K. Srivastava, J. Koutnik, B. R. Steunebrink, and J. Schmidhuber. 2017. "LSTM: A Search Space Odyssey," *IEEE Transactions on Neural Networks and Learning Systems*, 28(10):2222–2232. doi: 10.1109/TNNLS.2016.2582924.
16. A. Smola and B. Scholkopf. 2004. "A tutorial on support vector regression," *Statistics and Computing*, 14:199–222.
17. A. Ahmed and M. Khalid. 2018. "An intelligent framework for short-term multi-step wind speed forecasting based on Functional Networks," *Applied Energy*, 225(April):902–911. doi: 10.1016/j.apenergy.2018.04.101.
18. R. Zebari, A. Abdulazeez, D. Zeebaree, D. Zebari, and J. Saeed. 2020. "A Comprehensive Review of Dimensionality Reduction Techniques for Feature Selection and Feature Extraction," *Journal of Applied Science and Technology Trends*, 1(2):56–70. doi: 10.38094/jastt122.