



RNN-Based online anomaly detection in nuclear reactors for highly imbalanced datasets with uncertainty

Minhee Kim^a, Elisa Ou^b, Po-Ling Loh^c, Todd Allen^d, Robert Agasie^e, Kaibo Liu^{a,*}

^a Department of Industrial and Systems Engineering, University of Wisconsin–Madison, Madison, WI 53706, USA

^b Department of Electrical & Computer Engineering, University of Wisconsin–Madison, Madison, WI 53706, USA

^c Department of Statistics, University of Wisconsin–Madison, Madison, WI 53706, USA

^d Department of Nuclear Engineering and Radiological Sciences, University of Michigan, Ann Arbor, MI 48109, USA

^e Department of Engineering Physics, University of Wisconsin–Madison, Madison, WI 53706, USA

ABSTRACT

Accurate online condition monitoring and anomaly detection are crucial in nuclear applications to optimize economic performance and minimize safety risks. To achieve this goal, several major challenges exist which must be addressed. First, multi-sensor signals are often collected in the form of complex, multivariate time series. Second, relatively few anomaly records are available to train detection models. Lastly, the recorded data may contain uncertainties resulting from various sources, such as operator-induced variability and measurement error. In this paper, a recurrent neural network-based approach is proposed to tackle these issues by effectively utilizing historical data obtained during both normal and abnormal operations. Several advanced data preprocessing techniques are developed to improve the training process of the proposed neural network. The efficiency and sensitivity of the proposed method are evaluated on the multi-sensor signal measurements and operational reports obtained from a real case study. The results demonstrate much improved detection accuracy and practicality of the proposed method over conventional approaches.

1. Introduction

In nuclear power plants (NPPs), inexact condition monitoring and anomaly detection may result in unanticipated equipment failures and the associated economic losses, and even endanger public safety. To minimize such losses, a wide range of sensors that take measurements of temperature, pressure, neutron flux, etc., and other important system data, are installed throughout a nuclear reactor to monitor and ensure reliable, safe, and economical operation. In this paper, the term “anomaly” refers to acute anomalies, which is different from chronic anomalies and degradation processes that emerge slowly over time. Anomalies can be triggered by various causes affecting the relevant multi-sensor signals. Inference and assessment of the current plant status should then be made promptly and accurately, so that operators can carry out proper maintenance procedures. However, in a large NPP, it can be challenging for operators to process huge amounts of information of various formats and degrees of importance within a short period of time (Choi et al., 1995). Thus, it is highly desirable to develop a systematic process monitoring approach that can automatically assist operators in real-time anomaly detection. Additionally, as new designs are considered, these advances in monitoring approaches can be better integrated at the design phase.

In this study, we propose a novel anomaly detection method based

on artificial neural networks (ANNs) which fully utilizes the available data obtained during both normal and abnormal operations in NPPs. An ANN is a computational model that mimics the human brain learning process to learn complex non-linear relationships between inputs and outputs. Recently, ANN-based anomaly detection approaches have gained increasing attention in various applications due to their outstanding performance in dealing with complex systems (Khan and Yairi, 2018). In particular, we employ an recurrent neural network (RNN) structure, which is a special type of ANN that focuses on processing sequential data including sensor signals (Graves, 2012).

Developing an anomaly detection model based on multi-sensor signals in NPPs faces three major challenges: First, online multi-sensor signals are often available in the form of complex, multivariate time series, as different sensors measure various aspects of an NPP over fairly long periods of time. Accordingly, it is crucial to preprocess the raw data and systematically design the inputs and outputs of the model. Second, especially in nuclear applications, we are likely to have extremely imbalanced training datasets, since far more data are obtained during normal operations than abnormal operations. Third, the training dataset may involve uncertainties, especially when it contains manually recorded data. For instance, operational reports recording useful information such as the times associated with startup, scrams, and shutdown, are often available. Yet, as these reports are often manually

* Corresponding author at: Mechanical Engineering Building Room 3017, 1513 University Avenue, WI 53706-1572, USA.

E-mail address: kliu8@wisc.edu (K. Liu).

<https://doi.org/10.1016/j.nucengdes.2020.110699>

Received 15 January 2020; Received in revised form 27 April 2020; Accepted 11 May 2020

Available online 25 May 2020

0029-5493/ © 2020 Elsevier B.V. All rights reserved.

prepared, the information written on the reports may be inaccurate. Such uncertainties in nuclear datasets make anomaly detection far more challenging than typical anomaly detection tasks in other applications based on deterministic datasets. Here, we focus on label uncertainty, i.e., uncertainty in recorded anomaly times. When operators detect anomalies, they often either estimate anomaly times based on experience and domain knowledge, or record clock times after observing the anomalies and performing immediate follow-up actions. In both cases, exact anomaly times may be uncertain due to operator variability.

Our proposed method makes the following distinctive contributions, tailored to nuclear applications, by tackling the unique challenges mentioned above:

- We consider cases where historical anomalies are available and propose a novel RNN-based framework for anomaly detection in NPPs that effectively incorporates historical anomalies into model training. Here, different techniques are employed to address the major challenges resulting from the characteristic features of data collected in nuclear applications, including multi-sensor data, imbalance, and uncertainty.
- We propose a novel concept of robust labeling, which views labels probabilistically and can alleviate the issue of label uncertainty in anomaly detection.
- We conduct a series of numerical studies using the data collected from a real case study. The effects of the proposed techniques are investigated, and the anomaly detection performance is compared with a conventional residual-based method. The results show that the proposed method outperforms the conventional residual-based method and the improvement becomes more significant as the number of historical anomalies in the training dataset increases.

The rest of this paper is organized as follows. Section 2 reviews the existing approaches on online condition monitoring and anomaly detection in NPPs. Section 3 briefly introduces the concept of recurrent neural networks (RNNs) and principal component analysis (PCA). Section 4 describes the details of the proposed anomaly detection framework. Section 5 evaluates the accuracy and efficiency of the proposed method using data from a real case study. Section 6 presents conclusions and future research topics.

2. Literature review

In the literature, several efforts have been made to tackle online condition monitoring and anomaly detection in NPPs. In general, existing methods can be classified into two main categories: model-based and data-driven methods (Ma and Jiang, 2011). In model-based methods, a mathematical model is developed to represent the underlying physics of the system operation. Various model-based anomaly detection methods have been developed for nuclear applications, including Kalman filters (Roy et al., 1998; Tylee, 1983), parity equations (Gertler and Singer, 1990), and diagnostic observers (Frank et al., 1999; Frank and Ding, 1997). Yet, it may not be feasible to construct an accurate model for complex, nonlinear systems such as large NPPs, making the performance of existing algorithms often unsatisfactory.

Alternatively, data-driven methods build models based on a large amount of historical data with less required knowledge of the inherent

system operating physics. Popular data-driven methods for anomaly detection include singular value decomposition (SVD) (Mandal et al., 2017a), statistical process control (SPC) (Wang et al., 2018; Xian et al., 2018, 2019), support vector machines (SVMs) (Banerjee and Das, 2012; Zavaljevski and Gross, 2000), and artificial neural networks (ANNs) (Hadad et al., 2008; Mandal et al., 2017b; Messai et al., 2015). Recently, ANN-based approaches have been widely studied and shown outstanding performance especially on complex systems.

Most existing anomaly detection methods, including ANN-based approaches, train a model using only anomaly-free data obtained during normal operations to estimate the signal outputs. Anomalies are then detected when the residuals, i.e., the differences between true measurements and their estimates, exceed a predefined threshold (henceforth, we call these residual-based methods) (Hadad et al., 2008; Mandal et al., 2017a; Messai et al., 2015; Upadhyaya et al., 2003). Even when historical anomalies are available, they are only used to evaluate the constructed model and do not directly contribute to training the model. In fact, historical anomalies could be used in model training to learn unique patterns characterizing abnormal behaviors. It has been shown in different applications that incorporating both normal and abnormal training data may result in greatly improved anomaly detection accuracy, especially when future anomalies are expected to share similarities with historical anomalies (Gogoi et al., 2010).

To fill this literature gap, this paper aims at developing a novel RNN-based anomaly detection approach that systematically preprocesses the data obtained during both normal and abnormal operations and provides more accurate anomaly detection performance.

3. Introduction of recurrent neural networks and principal component analysis

This section contains two parts: In Section 3.1, we briefly introduce the concept of RNNs and review existing RNN-based anomaly detection approaches in nuclear applications. Combining RNNs with additional feature extraction methods can significantly improve training speed and model accuracy. In Section 3.2, we discuss the PCA algorithm which is one of the most widely used feature extraction methods and the benefits of using PCA with RNNs.

3.1. Recurrent neural networks

In complex systems such as nuclear reactors, it can be very time-consuming and expensive to manually extract high-level, meaningful features capturing complicated relationships between observed multi-sensor signals and underlying operational conditions. An ANN avoids this labor-intensive feature engineering step by directly learning features from the data itself. In particular, an ANN constructs models with layers which consist of nonlinear processing neurons. Each neuron learns to transform an input representation into an output representation at a more abstract level. By stacking multiple layers, we can build a network that learns a very complex function.

An RNN is a special type of ANN that focuses on processing sequential data such as speech, text, or time-series signals (Graves, 2012). Unlike conventional ANNs, an RNN has an internal loop and processes a sequence by iterating through the sequence elements and maintaining a “recurrent hidden state” that implicitly contains information about




Fig. 1. Illustration of a Recurrent Neural Network.

what it has seen so far, as illustrated in Fig. 1. A simple RNN, also known as a vanilla RNN, calculates a recurrent hidden state by

$$\mathcal{H}_t = \varphi(\mathbf{W} \cdot \mathbf{x}_t + \mathbf{U} \cdot \mathcal{H}_{t-1} + \mathbf{b}),$$

where $\mathbf{x}_t$ denotes an input at time step t (preprocessed multi-sensor signal measurements in our case), $\mathcal{H}_t$ denotes a hidden state vector at time step t , $\mathbf{W}$ denotes an input weight matrix, $\mathbf{U}$ denotes a hidden weight matrix, $\mathbf{b}$ denotes a vector of biases, and $\varphi(\cdot)$ is a nonlinear activation function. The hidden state at the last time step of the RNN is fed into another activation function $\varphi'(\cdot)$ to obtain the output $\hat{y}_t$ defined by

$$\hat{y}_t = \varphi'(\mathbf{V} \cdot \mathcal{H}_t),$$

where $\mathbf{V}$ is a weight matrix from the hidden layer to the output layer. Note that $\varphi(\cdot)$ and $\varphi'(\cdot)$ can be two different activation functions. For instance, for anomaly detection problems, sigmoid activation function $\varphi(x) = \frac{1}{1+e^{-x}}$ is often used as it scales $\mathbf{V} \cdot \mathcal{H}_t$ to the range 0 to 1 to indicate the normal and abnormal states of the system at the corresponding time. For a given training set $\{\mathbf{x}_t, y_t\}$, $t = 0, \dots, T$, the RNN aims at solving the following optimization problem for a given loss function $\mathcal{L}(\cdot, \cdot)$ via gradient descent to find the optimal weights and biases:

$$\min_{\mathbf{W}, \mathbf{U}, \mathbf{b}, \mathbf{V}} \mathcal{L}(y_t, \hat{y}_t).$$

Different variations of RNNs have been proposed including LSTMs and GRUs (Cho et al., 2014; Gers et al., 2000). Some of these variations have been applied in nuclear applications such as Infinite Impulse Response–Locally Recurrent Neural Networks (IIR–LRNNs) (Messai et al., 2015; Zio et al., 2009). However, we found that vanilla RNNs performed better than two of the most popular RNN variations, LSTM and GRU, in our case study (Section 5). One possible reason is that LSTM and GRU are quite computationally expensive to train and may cause overfitting and poor generalization on smaller dataset like in our case study. Thus, without loss of generality, we adopted the vanilla RNN structure in the paper. Note that extending the proposed method to other RNN variations is straightforward.

3.2. Principal component analysis

PCA is one of the most widely used methods for applications involving dimensionality reduction, data compression, and feature extraction (Jolliffe, 2011; Liu et al., 2013). Two popular definitions of PCA give rise to the same algorithm: First, PCA can be defined as a linear projection of the data onto a lower dimensional linear space, known as the principal subspace, which minimizes the average squared distance between the data points and their projections. Second, PCA can be defined as the orthogonal projection of the data such that the variance of the projected data is maximized.

Consider a set of N vectors $\{\mathbf{x}_n\}$, where $\mathbf{x}_n$ is D -dimensional, i.e., $\mathbf{x}_n = \{x_{n1}, x_{n2}, \dots, x_{nD}\}$. Using PCA, we find a linear projection of the D -dimensional vectors $\{\mathbf{x}_n\}$ onto a space with dimensionality $M \leq D$. It has been proved that the optimal linear projection in which the variance of the projected data is maximized (or the average squared distance between the original data and the projected data is minimized) occurs when the principal subspace is the span of the eigenvectors $\mathbf{u}_1, \dots, \mathbf{u}_M$ of the sample covariance matrix $\mathbf{S}$ corresponding to the M largest eigenvalues. The sample covariance matrix $\mathbf{S}$ is defined as

$$\mathbf{S} = \frac{1}{N-1} \sum_{n=1}^N (\mathbf{x}_n - \bar{\mathbf{x}})(\mathbf{x}_n - \bar{\mathbf{x}})^T,$$

where $\bar{\mathbf{x}}$ is the sample mean.

The optimal value of M is commonly decided by plotting the cumulative explained variance ratio $r(M)$ corresponding to the number of principal components as follows.

$$r(M) = \frac{\sum_{i=1}^M \lambda_i}{\sum_{j=1}^D \lambda_j},$$

where $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_D$ are the sorted eigenvalues of the sample covariance matrix $\mathbf{S}$. When $M = D$, this linear projection could be done with no information loss. $M < D$ leads to dimensionality reduction, data compression, and feature extraction. We can choose M to be the minimum value such that the cumulative explained variance ratio $r(M)$ is higher than a threshold value, e.g., 99% of the total variance. Then the projected data $\hat{\mathbf{x}}_n$ using the first M principal components could be obtained by $\hat{\mathbf{x}}_n = \mathbf{U}_M^T \mathbf{x}_n$, where $\mathbf{U}_M = [\mathbf{u}_1, \dots, \mathbf{u}_M]$ is the projection matrix.

While ANNs are able to learn complicated features from the raw data, combining them with additional feature extraction methods such as PCA can significantly improve model accuracy and speed up model training, especially when only a limited amount of training data is available. In our case, multi-sensor signals have high correlations (e.g., similar types of sensors or closely located sensors), so PCA can greatly reduce computational complexity and improve model accuracy without much information loss. As a result, we first apply PCA for dimension reduction and then use the compressed (projected) data to train the RNN. We will further examine how PCA can improve anomaly detection performance in Section 5.2.

4. Methodology

Section 4.1 first explains how we slide a fixed-width time window to generate multiple training samples of multi-sensor signals with the same length and then implement PCA to reduce the dimensionality of the data. Section 4.2 then describes how to label these training samples based on the manually recorded anomaly time to incorporate label uncertainty. As described in Section 1, it is common to have extremely imbalanced training samples in nuclear applications, i.e., far more training samples are obtained during normal operations than abnormal operations. In Section 4.3, we illustrate how to tackle this imbalanced training data via oversampling. Section 4.4 provides an overview of how to construct and train the proposed RNN-based framework based on the preprocessed dataset.

4.1. Sliding time window and PCA

Online multi-sensor signals in NPPs are often collected over long periods of time, and different sensors measure different aspects of an NPP simultaneously. To preprocess such data, one intuitive approach is to generate one training sample for each operation record. Yet, to capture the behavior of a system over a short time period and provide sufficient training samples, it is more efficient to generate multiple training samples from one operation record. We can achieve this by sliding a time window of fixed width n_{TW} over the whole multi-sensor signal data, i.e., dividing a long time series into a set of equal-sized subsequences. Fig. 2 illustrates how we slide a time window of width n_{TW} to generate multiple training samples from the data collected during one operation, when a system has d sensors. For any system whose operating time is shorter than n_{TW} , we can create one sample using its full operating time. The suitable width of the time window n_{TW} can be decided based on the characteristics of a system, e.g., signal sampling frequency or mechanisms of anomalies. We can either adopt domain knowledge or investigate historical signals (e.g., cross validation) to decide the appropriate value of n_{TW} . In general, while a larger n_{TW} conveys more information to the model, it also makes the model training procedure more computationally expensive.

After generating samples from the raw data, we apply a normalization technique to all signals, e.g., z-score normalization or min–max normalization, such that each signal is within a comparable range. Then the PCA algorithm introduced in Section 2.2 can be applied to reduce




Fig. 2. Sliding time window of width n_{TW} where a system has d sensors.

the dimensionality of the data.

4.2. Robust labeling

In this study, our primary goal is to accurately detect when anomalies occur. From a supervised learning perspective, this goal can be achieved by considering anomaly detection as a binary classification problem, where a label is 1 if the data is obtained when an anomaly

occurs, and 0 otherwise. For instance, suppose an anomaly time is recorded as T in the operational report. It is straightforward to assign label 1 to the samples (sub-sequences) collected at T , and 0 to the rest of the samples. However, as described in Section 1, the anomaly time recorded in an operational report may involve uncertainty, particularly when reports are manually recorded. In such cases, it is highly likely that the exact anomaly time is different from T .

To incorporate such type of uncertainty in the proposed framework, we develop a technique called “robust labeling.” Instead of labeling each data with binary codes, 0 (normal) and 1 (abnormal), robust labeling treats the probability that data is obtained when the abnormal event happens as the label, i.e., placing a distribution over the anomaly time. For instance, we can place a Gaussian distribution over the exact anomaly time with mean T and a certain variance; in this way, a larger probability value is assigned to the data collected closer to T . Alternatively, we can also use a uniform distribution with width $2\Delta t$, assigning probability $1/2\Delta t$ to the data collected between $T - \Delta t$ and $T + \Delta t$ and 0 to the rest. Domain knowledge or historical data can be used to select the appropriate distribution form and the parameters of the distribution, e.g., variance and Δt . In our case study, the recorded time on the operational report is the time when the operator completes immediate follow-up actions after detecting the anomaly, so we use a truncated Gaussian distribution with support $(-\infty, T]$, i.e., the probability that the anomaly happens after T is 0. This probability density function is then used to calculate the label of each sample. In particular, the label corresponding to the sample at time t can be interpreted as the probability that the abnormal event happens within the interval




Fig. 3. The proposed RNN-based online anomaly detection framework.




Fig. 4. UWNR open pool reactor.

($t - \delta t, t$), where δt is a very short time duration. The calculated labels are scaled such that the maximum possible label is equal to 1, since the ANN's outputs using the sigmoid activation function are within the range of [0, 1]. Note that a conventional method using binary codes can be regarded as a special case of the robust labeling approach, where we use the Dirac-delta distribution function, i.e., $y(t) = 1$ if $t = T$ and $y(t) = 0$ otherwise, where $y(t)$ is the label of the sample at time t . In some cases such as the Gaussian distribution, $y(t) > 0$ for all t . We then set $y(t)$ to be 0 if it is too small, assuming that the difference between the recorded abnormal time and the true abnormal time is not too large.

In fact, even when we have limited uncertainty regarding labels, e.g., the anomaly time is recorded with absolute certainty, robust labeling can still contribute to model accuracy and generalization. This is because we may view the robust labeling technique as capturing not only abnormal patterns in multi-sensor signals observed at the exact moment when an anomaly occurs, but also patterns observed right before or after the anomaly. In other words, robust labeling might be used to identify signs of a failure event before its occurrence, and thus it has the potential to be extended from detection of acute anomalies to detection of gradual/chronic anomalies or even anomaly prediction in cases when we have exact accurate anomaly times in the training data. However, we will leave this investigation to future work.

Table 1
Summary of scrams.

Scram Time	Description
2017/05/11 09:10:00	Reason for scram is unknown. No indication of trip element. Possible spurious failure of actual SCRAM relays. During completion of this step of the procedure an obvious power fluctuation occurred in the building and an actual Loss of Alternating current annunciator came in. Now suspect SCRAM is a result of building wide power fluctuation.
2017/12/01 11:21:00	SCRAM from picoammeter number 2. While performing a normal reactor setup, a reactor operator trainee did not appreciate the differential worth of the transient rod and inserted sufficient reactivity to result in a short period alarm. The trainee became distracted by the period alarm and failed to uprange the picoammeter to the next higher range.
2018/01/25 09:47:00	Manual emergency shutdown. During steady state operations at full power the ventilation system exhaust fan failed. The on-duty reactor operator observed decreasing exhaust flow and immediately initiated corrective actions by inserting a manual SCRAM.
2018/07/16 10:38:00	Shutdown by manual rundown following the loss of EF-7.

Table 2
Detailed description of sensor signals.

Symbol	Description	Units
LCR	Log Count Rate	Log CR
pA #1	Picoammeter #1	%
pA #2	Picoammeter #2	%
LogN	Log Neutron Power	Log W
Period	Period	DPM
Core Inlet	Core Inlet Temperature	°F
Demin Inlet	Demineralizer Inlet Temperature	°F
Pri HX I	Primary Heat Exchanger Inlet Temperature	°F
Pri HX O	Primary Heat Exchanger Outlet Temperature	°F
Int HX I	Intermediate Heat Exchanger Inlet Temperature	°F
Int HX O	Intermediate Heat Exchanger Outlet Temperature	°F
CWS	Chilled Water Supply Temperature	°F
CWR	Chilled Water Return Temperature	°F
Fuel	Fuel Temperature	°C
CAM Part Raw	Continuous Air Monitor Particulate Raw	
CAM Gas Raw	Continuous Air Monitor Gas Raw	
SAM Part Raw	Continuous Air Monitor Particulate Raw	
SAM Gas Raw	Continuous Air Monitor Gas Raw	
Vent Flow	Ventilation Flow Rate	KCFM
Demin Res	Demineralizer Resistivity	MOhm-cm
Hold Tank Level	Hold Tank Level	Gallon
Pool Level	Pool Level	Feet
CAM Part	Continuous Air Monitor Particulate Concentration	uCi/mL
CAM Gas	Continuous Air Monitor Gas Concentration	uCi/mL
SAM Part	Continuous Air Monitor Particulate Concentration	uCi/mL
SAM Gas	Continuous Air Monitor Gas Concentration	uCi/mL
CAM Part Exponent	CAM Part Exponent Concentration	uCi/mL
CAM Part Significand	CAM Part Significand	
CAM Gas Exponent	CAM Gas Exponent Concentration	uCi/mL
CAM Gas Significand	CAM Gas Significand	
SAM Part Exponent	SAM Part Exponent Concentration	uCi/mL
SAM Part Significand	SAM Part Significand	
SAM Gas Exponent	SAM Gas Exponent Concentration	uCi/mL
SAM Gas Significand	SAM Gas Significand	
Totalized Air Flow	Totalized Air Flow	MCF

4.3. Random oversampling

In practice, the number of operation records involving anomalies is much smaller than the number of records without anomalies. This data imbalance can lead to significant detection errors in most machine learning approaches, which are often designed to maximize accuracy on a balanced dataset (Buda et al., 2018; Yap et al., 2014). Two common ways to alleviate this problem are by oversampling data in the minority class or downsampling data in the majority class. In this study, we employ random oversampling, i.e., sampling with replacement from the available training samples representing anomalies. In this way, we balance the ratio between the number of samples representing anomalies and the number of samples obtained during normal operations. After applying the robust labeling method to the training data as




Fig. 5. An example of multi-sensor signals with a recorded anomaly marked by the vertical red line.

described in Section 4.2, samples with label > 0 are randomly over-sampled. In Section 5.2, we will investigate how much improvement in anomaly detection performance can be achieved using the random oversampling approach.

Random oversampling may lead to overfitting to the training data and poor generalization to the test dataset particularly when the future anomalies are significantly different from the historical anomalies. To avoid such overfitting issues, more advanced oversampling techniques have been developed, such as the Synthetic Minority Over-sampling TEchnique (SMOTE) and ADaptive SYNthetic Sampling (ADASYN) (Chawla et al., 2002; He et al., 2008). For instance, in SMOTE, new synthetic minority class samples (i.e., samples representing anomalies) are generated by randomly interpolating pairs of closest neighbors in the minority class. However, it has been shown that SMOTE does not work well for high-dimensional data such as the multi-sensor signals in this study (Blagus and Lusa, 2013). Thus, in our case study, we consider situations in which the anomalies in the test dataset share similarities

with those in the training dataset and apply random oversampling to handle imbalanced training data.

4.4. Overview of the proposed RNN structure

Fig. 3 illustrates the overall framework of the proposed approach. First, the data preprocessing procedures described in Sections 4.1 through 4.3—sliding time windows, PCA, robust labeling, and random oversampling—are conducted. The RNN then aims to learn relationships between the preprocessed multi-sensor signals and their corresponding labels. As the labels all lie in the range $[0, 1]$, we apply a sigmoid activation function and use binary cross-entropy as the loss function $\mathcal{L}(\cdot, \cdot)$ (Goodfellow et al., 2011). Depending on the characteristics of the data, e.g., the number of sensors, sampling frequency, and size of the training dataset, stacking multiple RNN layers may yield better performance by learning more complicated patterns hidden in the data.

UWNR 112 Rev 10 RSC Approval 5-18-16.docx Page 1 of 2

OPERATING LOG SHEET

TIME	CONTROL ELEMENT POSITION					PICOAMMETER		LOG N	FUEL TEMP INLET °C 19M1	CORE INLET TEMP °F	REMARKS	TS 6.1.3.B 2ND	ON DUTY RO	ON DUTY SRO
	No.1	No.2	No.3	Reg.	Trans	No.1(Range)	No.2(Range)							
0912						()	()	RO/SRO 2nd Person IN				AG	AM	SD
0916						()	()	Begin UWNR 110						
0938						()	()	Console ON						
0940						()	()	UWNR 110 Complete						
0910						()	()	SCRAM Occurred						
0937						()	()	RO / 2nd Change				AM	AG	
0949						()	()	Begin UWNR 115 Restart						
1001	11.54	11.54	11.54	11.54	11.54	100 (1 MW)	100 (1 MW)	100 W 301 77						
1100	11.57	11.57	11.57	11.57	11.57	100 (1 MW)	100 (1 MW)	100 W 303 75						
1101						()	()	RO / 2nd Change				AG	KZ	
1200	11.58	11.58	11.58	11.58	11.58	100 (1 MW)	100 (1 MW)	100 W 305 75						
1201						()	()	RO / 2nd Change				KZ	3Q	
1300	11.61	11.61	11.61	11.61	11.61	100 (1 MW)	100 (1 MW)	100 W 305.2 74						
1300						()	()	RO / 2nd Change				AG	KZ	
1356	11.63	11.63	11.63	11.63	11.63	98 (1 MW)	99 (1 MW)	100 W 303 74						
						()	()							
						()	()							
						()	()							
						()	()							
						()	()							

TIME CRITICAL 0906, 0954 TIME AT POWER 1000 POWER 1 MW TIME SHUT DOWN 0910, 1400

EXPERIMENTS IN PROGRESS Full Power Operations & Isotope Production

RUN NUMBER 1095 CORE CONFIGURATION K21-R6 DATE MAY 11 2017

Fig. 6. An example of an operating log (the recorded scram time is highlighted in the red bold rectangle).

5. Numerical study

In this section, we apply the proposed method to a real case study. In Section 5.1, we first provide an overview of the reactor and dataset. In Section 5.2, we apply different techniques developed in Sections 4.1 through 4.3 to the dataset and investigate how each of them can contribute to the anomaly detection performance. Section 5.3 demonstrates the benefits of using historical anomalies in model training. In particular, the proposed method is compared with a conventional residual-based method to investigate how anomaly detection performance improves according to the number of historical anomalies in the training data.

5.1. Overview of the reactor and dataset

The University of Wisconsin Nuclear Reactor (UWNR) is a heterogeneous, pool-type nuclear reactor currently fueled with low-enriched uranium TRIGA (Training, Research, Isotope Production, General Atomics) fuel and cooled by natural convection. Fig. 4 illustrates the UWNR open pool reactor. The reactor has been operating successfully for > 40 years, and supports a mission of education, training, and research. In this study, we use multi-sensor signals and operational reports collected over 40 operations from April 2015 to August 2018. Out of the 40 operations, 4 operations contained scrams. Recorded times and description logs of the scrams are summarized in Table 1. We consider these scrams as anomalies and investigate whether the proposed data-driven method can detect sudden changes and characterize distinct behaviors of multi-sensor signals due to scrams without prior information about scrams and the underlying mechanism of the reactor operation.

The data contains 35 sensors measuring different aspects of the system condition, i.e., $d = 35$. Table 2¹ provides detailed descriptions of these sensor signals, which are measured and recorded at one-second intervals. Fig. 5 shows an example of multi-sensor signals collected during an operation with an anomaly. In Fig. 5, red vertical lines indicate anomaly times recorded in operational reports. We can see that different sensors show significantly different patterns and trends. Moreover, a comparison between this operation record and other records shows that even the same sensor may behave quite differently under different operating circumstances. This complex data structure motivates us to explore the powerful ANN-based methods described earlier. The operational reports contain five different types: 3 checklists, 1 operating log, and 1 scram report. All these reports are handwritten. Figs. 6 and 7 show examples of an operating log and a scram report, respectively. As the focus of this study is on anomaly detection, we only use the two types of reports which describe the operating log and scram. In particular, the estimated anomaly time is obtained from the operating log, and the reason and follow-up action for this anomaly are obtained from the scram report. In cases where we have many handwritten documents, we may apply automatic hand-writing recognition algorithms to extract the information of interest. However, here, we only have 4 anomalies, so we can manually extract the information from the reports.

Note that for this dataset, the operator recorded the time after he/

¹ CR: Count Rate in units of counts per second, W: Watts (the power level of the reactor), DPM: Decades Per Minute (e.g., 1 DPM indicates that the reactor power goes from 10 Watts to 100 Watts in one minute), KCFM: Thousand Cubic Feet per Minute, and MCF: Mega Cubic Feet.

UWNR 115 Rev 6 RSC Approval 5-18-16 Page 1 of 2

SCRAM

The information required below will be recorded for each automatic or emergency manual scram or unintentional drop of one or more control elements unless scram was received when all control elements were in their full in position or drop was deliberate (rod drop reactivity measurement, etc.). A senior reactor operator (SRO) must be notified immediately upon occurrence of a scram. The reactor will not be operated until SRO permission is received by his or her signature on this form.

1. Type of scram. Automatic(✓) Emergency Manual()
 Unintentional drop of control element()
 Explain Unknown, no indication of trip element

2. Instrument channel/condition initiating scram unknown,
SCRAM, Blade Disengaged AM no indication of trip element

3. Annunciators giving an indication SCRAM, Blade Disengaged

4. Magnet currents after blade drop: #1 0 #2 0 #3 0

5. Reason for scram Unknown, no indication of trip element.
Possible spurious failure of actual SCRAM relays. During completion
of this step of the procedure an obvious power fluctuation occurred
in the building and an actual Loss of Alternating Current annunciator
came in. Now suspect SCRAM is a result of building wide power
fluctuation.

6. Follow-up action (Instrument repaired, etc.) None. Do not
suspect any instrument failure.

(Fuel temperature scram or power level above 1.25 MW steady state may require fuel element measurements-see T.S. 4.1.5.c.)

Fig. 7. An example of a scram report.

she observed the abnormal event and performed immediate follow-up actions. All the sensors and the clock in the control room acquire the time from the local area network (LAN), and are synchronized to the Network Time Protocol (NTP) server every hour. Yet, the control room clock does not display seconds (i.e., one-minute resolution). Depending on the complexity of the abnormal event and operator's expertise, the time difference between the exact abnormal event and the recorded time could be several seconds to several minutes.

5.2. Results of the proposed model

In this subsection, the anomaly detection performance of the

proposed model is investigated. We randomly sample five operations, two with anomalies and three without anomalies, to generate a test dataset, and then use the remaining 35 operations as the training dataset. A sliding time window with $n_{TW} = 600$ is applied to generate multiple samples of multi-sensor signals from each operation. The window width is determined using 10-fold cross validation to maximize the true positive ratio (TPR) = (# of true positives) / (# of true positives + # of false negatives), i.e., the proportion of anomalies that are correctly detected out of all anomalies. Here, we prioritize high TPR, since the failure of anomaly detection (false negatives) leads to much more deleterious results than false alarms (false positives) in nuclear applications. The following four models are considered to preprocess




Fig. 8. Cumulative explained variance ratio with respect to the number of components.

Table 3

The TPR, AUC, and average model training time of each model (the best performance is highlighted in bold and the standard deviation is in parentheses).

	TPR	AUC	Average Model Training Time (s)
Model (1)	0.669 (0.047)	0.744 (0.018)	237.93 (50.68)
Model (2)	0.820 (0.039)	0.809 (0.018)	204.05 (53.85)
Model (3)	0.866 (0.064)	0.827 (0.031)	215.74 (53.11)
Model (4)	0.922 (0.071)	0.849 (0.020)	336.82 (122.56)

Table 4

The AUC of each model (the best performance is highlighted in bold and the standard deviation is in parentheses).

		# of historical anomalies in training data	AUC
Proposed models	Model (A)	1	0.757 (0.126)
	Model (B)	2	0.821 (0.037)
	Model (C)	3	0.843 (0.019)
Benchmark model	Model (D)	0	0.691 (0.025)

the samples:

- (1) Proposed model not applying PCA, robust labeling (i.e., binary labeling of the data with 0 and 1), or oversampling
- (2) Proposed model applying PCA, but not robust labeling or oversampling
- (3) Proposed model applying PCA and robust labeling, but not oversampling
- (4) Proposed model applying PCA, robust labeling, and oversampling

From Model (1) to Model (4), we incrementally add techniques illustrated in Sections 4.1 to 4.3 to the model. Comparisons between different models highlight the benefit of each proposed data preprocessing procedure in accordance with the model construction. First, Fig. 8 plots the cumulative explained variance ratio, $r(M)$ as a function of the number of principal components M , using PCA. Consequently, we decided to use 21 principal components in Models (2) to (4), which explain 99.2% of the total variance. To extract eigenvalues and corresponding eigenvectors, the approximation approach proposed in (Halko et al., 2011) is used. Second, robust labeling is applied using a truncated Gaussian distribution in Models (3) to (4). This means that if the recorded time on the operating log is T , we use a Gaussian distribution with mean $T - 60$ and variance 300, and the distribution is truncated at

T (a label after T is set to 0). The variance is again chosen using 10-fold cross validation to maximize the TPR. In addition, if the label is less than 0.00001, it is set to 0. After generating samples from the raw multi-sensor signal data by using the sliding time window approach and applying robust labeling, the samples turn out to be highly imbalanced (i.e., the ratio between the number of samples with label > 0 and the number of samples with label $= 0$ is around 75). Lastly, in Model (4), the training samples with label > 0 are randomly oversampled to set the ratio between the number of samples with label > 0 and those with label $= 0$ to be 1.

Although Models (1) to (4) apply different techniques to preprocess the data and generate training samples, they all use the same structure of RNN to fairly compare the effects of the proposed techniques. We first fix the number of hidden layers to 1 to reduce computational costs and optimize the number of hidden neurons via 10-fold cross validation. As a result, the number of hidden neurons of the RNN is set to 50, which achieves the best overall performance. We also explored deeper structures of RNNs, i.e., RNNs with > 1 hidden layer, but the resulting anomaly detection performance was worse than the RNN with 1 hidden layer in our case study. One possible reason is that deeper networks are more likely to overfit the training data, as they involve a larger number of parameters.

RMS-Prop which is one of the most widely used optimization algorithms designed for NNs is used to train the RNN (Tieleman, 2012). RMS-prop adaptively sets the learning rate in a way that smaller learning rates are used for more frequently updated dimensions. The simulations are repeated 100 times for each model. Every trial was executed with two Intel(R) Xeon(R) CPU E5-4620 0 2.20 GHz processors and 192 GB RAM. All models are implemented and experimented using Tensorflow framework. Along with the TPR, the area under the receiver operating characteristic (ROC) curve, known as AUC, is employed to evaluate and compare the anomaly detection performance of different models. The AUC, which is one of the most commonly used metrics for classification performance, provides a relative tradeoff between the TPR and false positive ratio (FPR) = (# of false positives)/(# of true negatives + # of false positives). AUC takes a value between 0 and 1. A perfectly inaccurate anomaly detection model has an AUC of 0; a model that makes random guesses has an AUC of 0.5; and a perfectly accurate model has an AUC of 1, i.e., a higher AUC corresponds to better anomaly detection performance. The TPR, AUC, and average model training time of each model are summarized in Table 3. We can see that from Model (1) to Model (4), as we apply more proposed techniques, the models obtain higher TPR and AUC. Model (4), which applies all the proposed techniques illustrated in Section 4, achieves the highest TPR and AUC, thus giving the best anomaly detection performance. In addition, a comparison between Models (1) and (2) shows how PCA improves anomaly detection performance and speeds up model training. As Model (4) applies random oversampling to train the model, it takes longer for model training than other models. Yet, once the model is trained, all constructed models take a similar amount of time to implement and detect anomalies in real time, as they have similar NN structures.

5.3. Comparison with the Residual-Based model

In this subsection, the anomaly detection performance of the proposed method is further compared with a conventional residual-based ANN method. As described in Section 1, most existing ANN-based approaches only use anomaly-free data from normal operations. An anomaly is identified when the residual (the difference between the observed and estimated values) is significant. Here, we construct a residual-based benchmark model which is also based on an RNN. To achieve a fair comparison, this benchmark model uses the same structure of RNN as the proposed model, i.e., 1 hidden layer with 50 hidden neurons, and the same data preprocessing procedures, including sliding time windows and PCA. The difference between the proposed model

and the benchmark model is that the benchmark model is trained only using anomaly-free data, whereas the proposed model is trained using the proposed robust labeling technique and data from both normal and abnormal operations. In the literature, different methods have been proposed to choose the optimal threshold value, but a standardized procedure does not exist. Thus, we use AUC, which provides a comprehensive measure of anomaly detection performance across all possible threshold values, to evaluate different models. Similar to Section 5.2, we consider the following models:

- (A) Proposed model when one operation with a historical anomaly is included in the training dataset
- (B) Proposed model when two operations with historical anomalies are included in the training dataset
- (C) Proposed model when three operations with historical anomalies are included in the training dataset
- (D) Benchmark model which aims to estimate the signal measurements and detects anomalies based on the residuals

We randomly sample four operations, one with anomalies and three without anomalies, to generate the test dataset. Then, 34 operations are randomly sampled from the remaining 36 operations to use as the training dataset, so that n_a operations have historical anomalies and $34 - n_a$ operations do not have historical anomalies, where n_a is 1, 2, 3, and 0 for Models (A), (B), (C), and (D), respectively. As a result, all the models use the same size of the training dataset, while only the number of historical anomalies in the training dataset varies. Other detailed settings of the proposed models and preprocessing procedures follow Section 5.2. The AUC of each model is summarized in Table 4. We can see that as the number of historical anomalies in the training dataset increases, the AUC of the proposed model increases. This is because as the proportion of historical anomalies in the training dataset increases, the RNN is exposed to more diverse realizations of multi-sensor signals during training, and thus it generalizes better in the testing phase compared with that of fewer anomalies (Chawla, 2005). Furthermore, the standard deviation of the AUC decreases with the number of historical anomalies, implying more stable anomaly detection performance. Yet, a comparison between Models (A) and (D) shows that even only one historical anomaly in the training data can provide valuable guidance to the model and significantly improve anomaly detection performance.

6. Conclusion

In this study, we proposed a novel RNN-based method based on several advanced techniques to tackle the practical challenges involved in analyzing nuclear data. RNNs have shown outstanding performance and potential in dealing with a wide range of complex systems, including nuclear applications. Nevertheless, most existing RNN-based anomaly detection methods only employ anomaly-free data during model training. Moreover, even when historical anomalies are available, there are still several challenges in incorporating the data into model training. To address the challenges, PCA and sliding time windows are applied to preprocess the raw data to capture the unique patterns of observed signals over a period of time, speed up model training, and minimize overfitting. Robust labeling is then developed to incorporate uncertainties involved in estimating historical anomaly times. Finally, the preprocessed training samples representing anomalies are randomly oversampled to balance the ratio between the numbers of anomalous samples and anomaly-free samples.

Comprehensive numerical studies on the real reactor dataset showed how each proposed technique contributes to anomaly detection performance. The proposed method outperforms the benchmark method trained only using anomaly-free data, measured in terms of AUC. The results demonstrate the importance of incorporating historical abnormal data into model training and the feasibility of using ANN-

based approaches for anomaly detection in NPPs.

Several topics exist for future work. First, it would be interesting to extend the proposed method to handle anomaly classification. For example, in cases where the types of historical anomalies are known, we can construct a similar RNN to the proposed one, designing the outputs with one-hot encoding and using the softmax activation function. Second, the proposed method considers the cases where we only know approximate anomaly times. It is worth exploring cases where we know exact anomaly times and extend the proposed model to predict future anomalies as well.

CRedit authorship contribution statement

Minhee Kim: Conceptualization, Methodology, Software, Writing - original draft. **Elisa Ou:** Software, Validation, Writing - review & editing. **Po-Ling Loh:** Writing - review & editing. **Todd Allen:** Writing - review & editing. **Robert Agasie:** Resources, Data curation, Writing - review & editing. **Kaibo Liu:** Conceptualization, Writing - review & editing, Supervision, Project administration, Funding acquisition.

Acknowledgements

The authors acknowledge the funding support in part by the office of Naval Research under Grant N00014-17-1-2261 and in part by the Department of Energy under award number DE-NE0008805.

References

- Banerjee, T.P., Das, S., 2012. Multi-sensor data fusion using support vector machine for motor fault detection. *Inf. Sci. (N.Y.)* 217, 96–107. <https://doi.org/10.1016/j.ins.2012.06.016>.
- Blagus, R., Lusa, L., 2013. SMOTE for high-dimensional class-imbalanced data. *BMC Bioinformatics* 14, 106. <https://doi.org/10.1186/1471-2105-14-106>.
- Buda, M., Maki, A., Mazurowski, M.A., 2018. A systematic study of the class imbalance problem in convolutional neural networks. *Neural Networks* 106, 249–259. <https://doi.org/10.1016/j.neunet.2018.07.011>.
- Chawla, N.V., 2006. Data Mining for Imbalanced Datasets: An Overview, in: *Data Mining and Knowledge Discovery Handbook*. Springer-Verlag, pp. 853–867. https://doi.org/10.1007/0-387-25465-x_40.
- Chawla, N.V., Bowyer, K.W., Hall, L.O., Kegelmeyer, W.P., 2002. SMOTE: Synthetic minority over-sampling technique. *J. Artif. Intell. Res.* 16, 321–357. <https://doi.org/10.1613/jair.953>.
- Cho, K., Van Merriënboer, B., Gulcehre, C., Bahdanau, D., Bougares, F., Schwenk, H., Bengio, Y., 2014. Learning Phrase Representations using RNN Encoder-Decoder for Statistical Machine Translation, in: *Proc. Conference on Empirical Methods in Natural Language Processing*. pp. 1724–1734.
- Frank, P.M., Ding, X., 1997. Survey of robust residual generation and evaluation methods in observer-based fault detection systems. *J. Process Control* 7, 403–424. [https://doi.org/10.1016/S0959-1524\(97\)00016-4](https://doi.org/10.1016/S0959-1524(97)00016-4).
- Frank, P.M., Schrier, G., García, E.A., 2007. Nonlinear observers for fault detection and isolation, in: *New Directions in Nonlinear Observer Design*. Springer London, pp. 399–422. <https://doi.org/10.1007/bfb0109937>.
- Gers, F.A., Schmidhuber, J., Cummins, F., 2000. Learning to Forget: Continual Prediction with LSTM. *Neural Comput.* 12, 2451–2471. <https://doi.org/10.1162/089976600300015015>.
- Gertler, J., Singer, D., 1990. A new structural framework for parity equation-based failure detection and isolation. *Automatica* 26, 381–388. [https://doi.org/10.1016/0005-1098\(90\)90133-3](https://doi.org/10.1016/0005-1098(90)90133-3).
- Gogoi, P., Borah, B., Bhattacharyya, D.K., 2010. Anomaly detection analysis of intrusion data using supervised & unsupervised approach. *J. Conver. Inf. Technol.* 5, 95–110. <https://doi.org/10.4156/jcit.vol5.issue1.11>.
- Goodfellow, I., Bengio, Y., Courville, A., 2011. *Deep learning*. MIT Press, The.
- Graves, A., 2012. *Supervised sequence labelling with recurrent neural networks*. Springer.
- Hadad, K., Mortazavi, M., Mastali, M., Safavi, A.A., 2008. Enhanced neural network based fault detection of a VVER nuclear power plant with the aid of principal component analysis. *IEEE Trans. Nucl. Sci.* 55, 3611–3619. <https://doi.org/10.1109/TNS.2008.2006491>.
- Halko, N., Martinsson, P.G., Tropp, J.A., 2011. Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions. *SIAM Rev.* 53, 217–288. <https://doi.org/10.1137/090771806>.
- He, H., Bai, Y., Garcia, E.A., Li, S., 2008. ADASYN: Adaptive synthetic sampling approach for imbalanced learning. In: *Proceedings of the International Joint Conference on Neural Networks*, pp. 1322–1328. <https://doi.org/10.1109/IJCNN.2008.4633969>.
- Jolliffe, I., 2011. *Principal Component Analysis*, in: *International Encyclopedia of Statistical Science*. Springer Berlin Heidelberg, Berlin, Heidelberg, pp. 1094–1096. https://doi.org/10.1007/978-3-642-04898-2_455.
- Khan, S., Yairi, T., 2018. A review on the application of deep learning in system health

- management. *Mech. Syst. Signal Process.* 107, 241–265. <https://doi.org/10.1016/j.ymssp.2017.11.024>.
- Liu, K., Jain, S., Shi, J., 2013. Physician performance assessment using a composite quality index. *Stat. Med.* 32, 2661–2680. <https://doi.org/10.1002/sim.5710>.
- Ma, J., Jiang, J., 2011. Applications of fault detection and diagnosis methods in nuclear power plants: A review. *Prog. Nucl. Energy* 53, 255–266. <https://doi.org/10.1016/j.pnucene.2010.12.001>.
- Mandal, S., Santhi, B., Sridhar, S., Vinolia, K., Swaminathan, P., 2017a. Sensor fault detection in Nuclear Power Plant using statistical methods. *Nucl. Eng. Des.* 324, 103–110. <https://doi.org/10.1016/j.nucengdes.2017.08.028>.
- Mandal, S., Santhi, B., Sridhar, S., Vinolia, K., Swaminathan, P., 2017b. Nuclear Power Plant Thermocouple Sensor Fault Detection and Classification using Deep Learning and Generalized Likelihood Ratio Test. *IEEE Trans. Nucl. Sci.* 1–1. <https://doi.org/10.1109/TNS.2017.2697919>.
- Messai, A., Mellit, A., Abdellani, I., Massi Pavan, A., 2015. On-line fault detection of a fuel rod temperature measurement sensor in a nuclear reactor core using ANNs. *Prog. Nucl. Energy* 79, 8–21. <https://doi.org/10.1016/j.pnucene.2014.10.013>.
- Roy, K., Banavar, R.N., Thangasamy, S., 1998. Application of fault detection and identification (fdi) techniques in power regulating systems of nuclear reactors. *IEEE Trans. Nucl. Sci.* 45, 3184–3201. <https://doi.org/10.1109/23.736198>.
- Choi, Seong Soo, Kang, Ki Sig, Kim, Han Gon, Chang, Soon Heung, 1995. Development of an on-line fuzzy expert system for integrated alarm processing in nuclear power plants. *IEEE Trans. Nucl. Sci.* 42, 1406–1418. <https://doi.org/10.1109/23.467727>.
- Tieleman, T., Hinton, G., 2012. RmsProp: Divide the gradient by a running average of its recent magnitude, in: COURSERA: Neural Networks for Machine Learning.
- Tylee, J.L., 1983. On-Line Failure Detection in Nuclear Power Plant Instrumentation. *IEEE Trans. Automat. Contr.* 28, 406–415. <https://doi.org/10.1109/TAC.1983.1103240>.
- Upadhyaya, B.R., Zhao, K., Lu, B., 2003. Fault monitoring of nuclear power plant sensors and field devices. *Prog. Nucl. Energy* 43, 337–342. [https://doi.org/10.1016/S0149-1970\(03\)00046-5](https://doi.org/10.1016/S0149-1970(03)00046-5).
- Wang, A., Xian, X., Tsung, F., Liu, K., 2018. A spatial-adaptive sampling procedure for online monitoring of big data streams. *J. Qual. Technol.* 50, 329–343. <https://doi.org/10.1080/00224065.2018.1507560>.
- Xian, X., Li, J., Liu, K., 2019. Causation-Based Monitoring and Diagnosis for Multivariate Categorical Processes with Ordinal Information. *IEEE Trans. Autom. Sci. Eng.* 16, 886–897. <https://doi.org/10.1109/TASE.2018.2873365>.
- Xian, X., Wang, A., Liu, K., 2018. A nonparametric adaptive sampling strategy for online monitoring of big data streams. *Technometrics* 60, 14–25. <https://doi.org/10.1109/COASE.2017.8256208>.
- Yap, B.W., Rani, K.A., Abd Rahman, H.A., Fong, S., Khairudin, Z., Abdullah, N.N., 2014. An application of oversampling, undersampling, bagging and boosting in handling imbalanced datasets. Springer Verlag, pp. 13–22.
- Zavaljevski, N., Gross, K.C., 2000. Sensor fault detection in nuclear power plants using multivariate state estimation technique and support vector. VINCA, Yugoslavia.
- Zio, E., Broggi, M., Pedroni, N., 2009. Nuclear reactor dynamics on-line estimation by Locally Recurrent Neural Networks. *Prog. Nucl. Energy* 51, 573–581. <https://doi.org/10.1016/j.pnucene.2008.11.006>.