INL/CON-24-77042-Revision-0

Statistical and Neural Network for Real Sensor-Data-Driven Anomaly Detection in Nuclear Applications

April 2024

Changing the World's Energy Future

Imtiaz Nasim, Vivek Agarwal, Joshua E Daw, Mingyue Ji, Sneha Kumar Kasera

INL/CON-24-77042-Revision-0

Statistical and Neural Network for Real Sensor-Data-Driven Anomaly Detection in Nuclear Applications

Imtiaz Nasim, Vivek Agarwal, Joshua E Daw, Mingyue Ji, Sneha Kumar Kasera

April 2024

Idaho National Laboratory
Idaho Falls, Idaho 83415

http://www.inl.gov

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

Statistical and Neural Network for Real Sensor-Data-Driven Anomaly Detection in Nuclear Applications

Imtiaz Nasim[1], Vivek Agarwal[1], Joshua E. Daw[1], Mingyue Ji[2], and Sneha Kasera Kumar[3]

Idaho National Laboratory, 1955 Fremont Avenue, Idaho Falls, Idaho 83415[1]
Department of Electrical & Computer Engineering, University of Utah[2]
Kahlert School of Computing, University of Utah[3]

[leave space for DOI, which will be inserted by ANS – do not highlight this line]

ABSTRACT

Anomaly detection (AD) in sensor data is critical to ensure uninterrupted functionality of nuclear power plants (NPPs). Consequently, AD model validation through real-world sensor data is important for applications in nuclear facilities. In this paper, we propose an Autoencoder (AE)—a multi-layered neural network, for AD in sensor data from an operational NPP testbed. Since the dataset lacks labels for irregularities, we introduce random noise and label them to effectively train our model. The proposed AE model assigns a higher reconstruction error to the abnormal samples that deviate from those encountered during the training phase and uses the reconstruction loss to detect anomalies in a *representative imbalanced dataset*. We also introduce an analytical solution—seasonal trend decomposition (STD)—as another AD scheme for identifying irregularities within the same time-series dataset. In contrast to the AE model which relies on reconstruction loss, the STD scheme decomposes the entire dataset into its trend, seasonality, and residual components to pinpoint irregularities. Our findings indicate that the proposed AE and STD models individually achieve recall scores of 97% and 92%, respectively. We validate the performance of the two models on both balanced and imbalanced data. We further solidify the results by picking the combined selected anomalies of the two solutions with an "AND" operator for more reliable predictions.

Keywords: Anomaly Detection, Seasonal Trend Decomposition, Autoencoder, Sensors, Nuclear Plant

1. INTRODUCTION

Anomaly detection (AD) techniques are used to identify instances that do not follow the regular behavior of the normal operating state of the system. Anomalies are deemed significant as they display rare events that deviate from normal operation, prompting critical actions across various application domains [1]. These domains encompass image processing, autonomous driving, fraud detection, fault identification, and abnormalities in industrial machines, among others. Nuclear power plants (NPPs) heavily depend on sensor-data-driven AD techniques to avoid failure and to ensure continuity of service [2]. Specifically, AD is crucial for advanced NPPs where condition-based monitoring and situational awareness are key to maintaining uninterrupted, failure-free operations.

Numerous machine learning (ML) models have been proposed in the literature to discover unique and rare patterns of outliers within the data [3]. For instance, support vector machine, which is a supervised ML technique with labeled training, was employed for classification-based AD tasks [4]. Such models have also been proposed as unsupervised solutions with adjustments in the training data features [5]. On the other hand, studies on clustering-based AD techniques assume that anomaly samples do not conform to the cluster of normal samples. Any subsequent new data that do not fit well within the existing cluster are considered anomalies [6]. Distance score was used as a metric to detect such anomalous events [7]. However, these models perform well on a balanced dataset, but they are not always suitable to detect anomalies in many real-world scenarios when the data are *imbalanced*.

In addition to classification and clustering, the other ML solutions are mostly residual-based ones. Such algorithms use clean samples for training and assign anomaly scores to the unseen testing samples based on their dissimilarity with the learned normal paradigm. Among the residual-based solutions, Autoencoder (AE), which is a neural network model, has gained significant interest in literature [8], [9] and was applied successfully over balanced and imbalanced data [10], [11]. The main principle of AE is to learn the compressed representation of raw data using dimensionality reduction and then use the reconstruction error for detecting outliers. The authors in [8] utilized the reconstruction probability of variational AE, while [9] considered the non-linear correlations between features to detect anomalies. Unlike principal component analysis (PCA), AE is a more robust unsupervised technique to treat the non-linearity of data [12] due to the AE's powerful generalization capability. Such data non-linearity is common in sensor-data-driven samples from NPPs. However, most so-called unsupervised AD techniques in literature assume that anomaly-free data are available for training while data under abnormal conditions are rare to achieve. This is not always true in real-world settings, where the training data may also be contaminated with a small fraction of abnormal instances. It is more appropriate to refer to such solutions as "semi-supervised" instead of fully unsupervised since they assume anomaly-free training data [13].

In addition to ML models, statistical AD solutions have also been proposed. Statistical methods attempt to fit the distributions on the training data, then any data with low probability under this distribution are treated as anomalous or outliers. Among the statistical models, the authors in [14] studied a Markov model-based AD approach which was validated with limited real data. Intrusion detection using AD was studied using the chi-square theory in [15]. The work in [16] proposed a statistical processing unit to detect anomalous network traffic. A statistical signal processing technique based on abrupt change in the signal was presented in [17]. However, such solutions were considered mostly for balanced datasets and were confined to artificially generated data without being widely applied on real data, except in some few studies such as [18]. Consequently, such solutions have rarely been considered for AD in *real sensor data from NPPs*, as in this paper. In summary, while there are schemes that leverage novel ML techniques [4]–[9] or statistical models [14]–[17] for AD, they mostly focus on artificially generated data or have limited validation on imbalanced datasets.

In this paper, we first propose a neural network-based AE model for detecting anomalies in sensor data from a nuclear testbed. For optimal performance, the AE requires the 'normal' (i.e., error-free) samples only for training so that it can yield high reconstruction loss for abnormal instances during the testing phase. But the data from real nuclear sensors are unlabeled, making it cumbersome to distinguish the normal and abnormal samples in advance. To address this challenge and toward solving a real-world problem with imbalanced data, we intentionally introduce random impulse noise to uniformly selected *1% samples* and labeled them as anomalous. Then, we use only the normal samples for training the AE, while the rest of the information is used for testing and validation. As a complementary approach, we also applied another solution, namely seasonal trend decomposition (STD), that detects abnormal instances by decomposing the sensor dataset into its trend, seasonality, and residual components. In contrast to our multilayered neural network-based AE solution, the proposed STD model adopts a statistical approach, detecting anomalies through the residual discrepancy between predicted and original values. The proposed STD approach learns the trend and seasonality of the given time-series data and uses a threshold to detect any dissimilarity in the new samples. For the sake of fair comparison between the two proposed strategies, we apply the same noisy imbalanced dataset for STD performance evaluation. The contributions of this paper are highlighted as follows:

- We formulate the AD problem as an AE problem where the reconstruction loss between the normal and anomaly samples are used to detect anomalies. We inserted a small noise fraction to the unlabeled dataset to make it labeled. This facilitates successful training of the model with a representative imbalanced data. The solution is applied over a real-world sensor dataset derived from an NPP.

- We also develop a statistical STD solution that decomposes the dataset into underlying seasonality, trend, and residual components to detect the anomalous data samples.
- Our models are validated across different types of datasets. For imbalanced data, the proposed AE model demonstrates a significant advantage in anomaly detection, while in scenarios with balanced data, the statistical model is sufficient.
- To enhance the robustness of our approach, we further combine the outputs of the two strategies. The combined output selects a sample as anomalous only when both solutions claim it abnormal.

2. DATA DESCRIPTION AND PROBLEM STATEMENTNT

This section first provides a brief explanation of the input samples used for our AD models. Then, for detecting the anomalies in real data, we derive the problem formulation of this paper.

2.1. Data Description

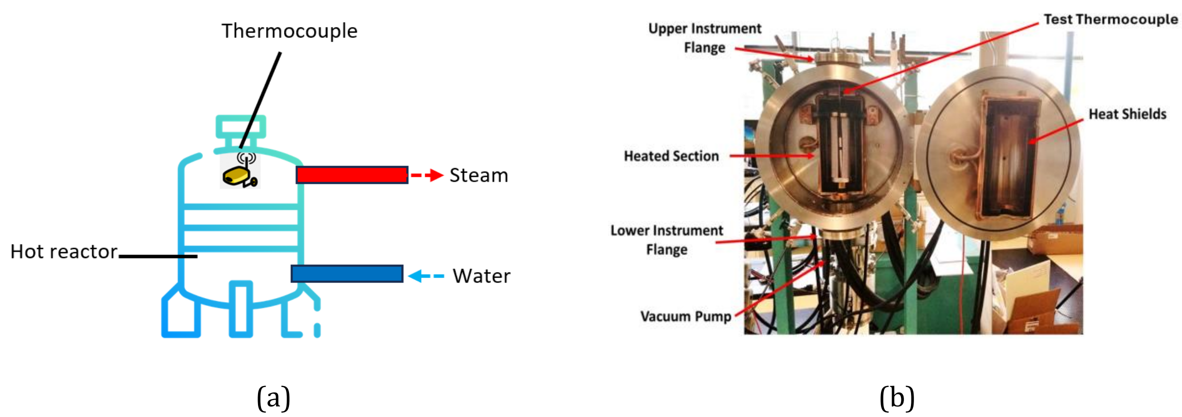

Figure 1. (a) Schematic of thermocouple sensor on a hot nuclear surface; (b) Actual vacuum furnace used for the test.

We use the output of thermocouple sensors that measure the temperature of a hot furnace in an NPP testbed at Idaho National Laboratory. For a change in temperature, an associated change in voltage is detected and measured by the thermocouples. We consider these voltage measurements as the input to our AD models. Due to the nature of nuclear sensors, the data acquired from thermocouples are unlabeled. This makes it challenging to select the instances in advance that exhibit irregular non-linear patterns that deviate from the regular regime. To address this problem and to meet the AE model's training requirement, we insert impulse noise to uniformly selected 1% of the entire dataset comprising 12,188 values. The impulse noise is applied by increasing the signal magnitude by a random number up to 30%. Samples with noisy signals exceeding a specific threshold are identified as anomalies and stored as ground-truths for later comparison. This method of adding noise to a minority class of samples enabled us to create a representative imbalanced dataset.

2.2. Problem Formulation

Let us assume that after insertion of noise, the training set becomes $X = \{x_1, x_2, \ldots, x_n\}$, and each of the elements is a d dimensional vector ($x_i \in R^d$). Using the samples in the training phase, we want to reproduce the data to get the output $X' = \{x'_1, x'_2, \ldots, x'_n\}$. We optimize the model to minimize the reconstruction error to get the optimal subspace. The error or residual in the reconstruction phase is defined as:

$$\varepsilon(x_i, x'_i) = \sum_{i=1}^{n} (x_i - x'_i)^2 \tag{1}$$

Consequently, the goal is to minimize the average error with respect to x_i and x'_i

$$x_i\ ,\ x'_i = \underset{x_i, x'_i}{arg\,min}\ 1/n \sum_{i=1}^{n} \varepsilon(x_i - x'_i)^2 \tag{2}$$

We formulate our anomaly detection problem as finding the samples $x_i \in X$ which produces a residual that exceeds a threshold after reconstruction using the following formula,

$$\begin{aligned} &\text{find } x_i \\ &\text{s.t } \varepsilon(x_i, x'_i) > \theta \end{aligned} \tag{3}$$

where θ is the chosen threshold value. Consequently, the reconstruction residual of a sample exceeding this threshold is considered as an anomaly.

3. PROPOSED ANOMALY DETECTION MODEL

This section further explains our AD techniques. First, we explain our AE model, which is a popular neural network developed on an encoding and decoding principle. Second, we provide details on the statistics-based AD technique which uses the seasonality and trend of the time series data to detect anomalies.

3.1. Autoencoder for Anomaly Detection

Autoencoder (AE) is a neural network that uses an encoder and a decoder to reconstruct the signal. The proposed AE solution is a *dimensionality reduction* scheme that utilizes reconstruction errors to detect anomalies. We first pass the inputs of the algorithm, which includes the noisy data samples, to an encoder that compresses the data into a latent space. The latent space is referred to as the hidden layer. Such hidden layer, sometimes also referred to as "code," is a low-dimensional and non-linear representation of the input [19]. The compressed input from the hidden layer is finally fed to the decoder, which decodes the data back to its original form with the same dimension. More specifically, the decoder takes the data from the hidden layer to the reconstruction phase where the dimensionality of the output is equal to the input. During the reconstruction phase, the AE model aims to get as close as possible to the input. But for an anomalous event which was not encountered during the training (due to training with normal samples only), it provides a larger reconstruction error, thus deviating from the clean data.

AE learns a mapping solution from the input to itself through a pair of encoding and decoding phases,

$$X' = D(E(X)) \tag{3}$$

where X is the input data, E is the encoding map from the input data to the hidden layer, D is a decoding map from the hidden layer to the output layer, and X' is the recovered version of the input data. The goal is to train the encoder E and decoder D so that the difference between X and X' is minimized. In particular, an AE is a solution to the following optimization problem [19]:

$$\min_{D,E}\ ||X - D(E(X))|| \tag{4}$$

where ||.|| is chosen to be the l_2 norm [19].

Note that unlike any supervised learning that uses labeled samples, the AE training phase is unsupervised as it does not require the labels of the training samples. Instead, the training requires the features of the normal samples only so that the model can produce a larger reconstruction loss for the abnormal samples. Usually, an AE with more than one hidden layer is called a deep autoencoder [20] and each additional hidden layer requires an additional pair of encoders $E(\cdot)$ and decoders $D(\cdot)$. However, for simplicity in this paper, we used only one hidden layer for our AE model to represent the reduced input dimension.

3.2. Anomaly Detection with Seasonal Trend Decomposition

Our statistical model is an AD scheme that decomposes the time series data into its seasonality, trend, and residual components toward identifying the anomalous points. Once the trend and seasonal components are separated, the remaining residuals are used to identify the outliers or anomalies. For the given data X, assume its trend component is X_T, and the seasonal and residual components are X_S *and* X_R, respectively. Then the original signal can be denoted as,

$$X = X_T + X_S + X_R \tag{5}$$

The trend is used to capture the nature of increasing or decreasing direction within the data. The seasonality captures the seasonal factors over a fixed period. When the signal is decomposed to its trend and seasonality, the leftovers are the residuals that indicate the noise in the given time-series data [21]. Therefore, from equation (1) the residual is simply calculated as,

$$X_R = \mathrm{X} - (X_S + X_T). \tag{6}$$

Our STD model uses LOESS (Local RegrESSion) smoothing, which is one of the most popular STD algorithms (are often termed as STL instead [22]). As such, and to be consistent with previous studies, we will address our STD algorithm as STL for the rest of this paper, unless otherwise mentioned. We assume a confidence interval of two standard deviations in the distribution of errors as the threshold θ for our STL model. Any point residing outside θ is considered a potential anomaly in the given time-series data.

4. RESULTS

We used PYTHON for implementation of the proposed models. ReLU wass used as the activation function for AE neural network.

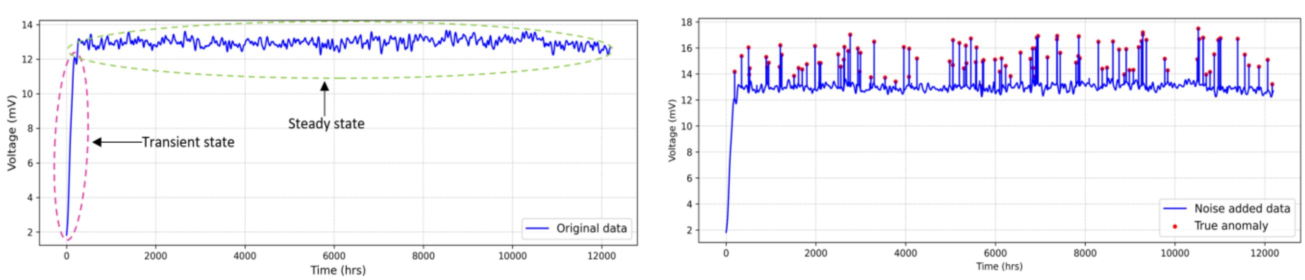

(a) (b)

Figure 2. (a) Original data; (b) data after 1% noise insertion.

Figure 2(a) illustrates the original data collected from real nuclear thermocouple sensors, measuring voltage in millivolts (mV) in response to temperature changes. The transient state of the data lasts for 1.5 hours (which is less than 2% of the experimentation time) before it reaches the steady state. The scope of this paper is not to detect anomalies in the transient condition but rather to detect them in the steady state, which mostly informs NPP operational decisions. Figure 2(b) shows the dataset after introducing noise to 1% of uniformly selected samples. Instances where the noise increases the signal magnitude by more than 0.5 of its previous state are labeled as anomalous (label 1), while the rest are labeled as clean or normal (label 0).

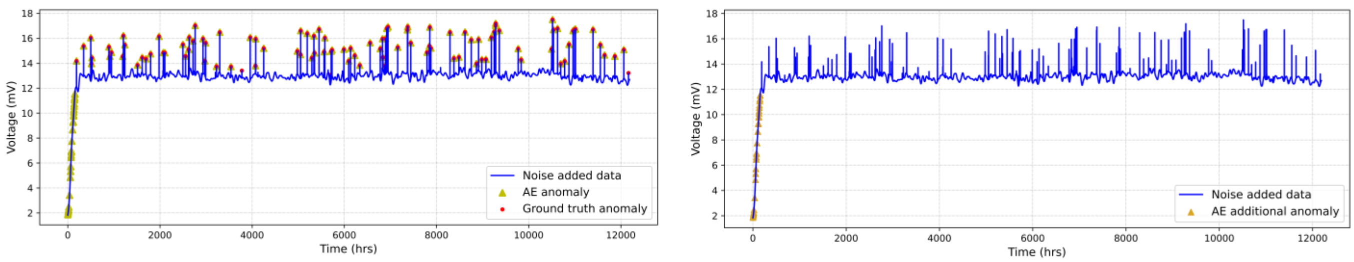

(a) (b)

Figure 3. (a) AE detected anomalies; (b) AE false positives.

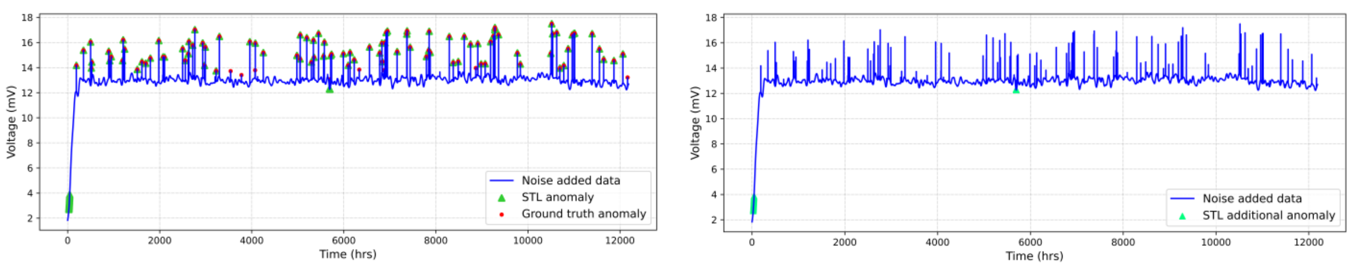

(a) (b)

Figure 4. (a) STL detected anomalies; (b) STL false positives.

Ground-truth anomalous labels, depicted as red dots in Figure 2(b), identify only 101 out of the total 12,188 instances. This imbalance in the data with anomalies representing a minority portion often reflects real-world scenarios in data collection.

4.1. Performance Evaluation of AE and STL model

The comparison between the AE model's predicted anomalies and ground-truth anomalies is depicted in Figure 3(a). The proposed AE model successfully tracks 97% of the ground-truth anomalies. However, the model also predicts some clean samples as anomalies, leading to false positives as depicted in Figure 3(b). Such false alarms are confined to the transient state only and disappear once the signal reaches its steady state after just 140 samples.

Figure 4 shows a similar performance of the statistical STL model. We consider the same imbalanced noisy dataset for STL to ensure fairness in comparison. Similar to AE, the proposed STL model detects more than 92% of the actual anomalies as shown in Figure 4(a), and incorrectly detects some clean instances as anomalous, as shown in Figure 4(b). Unlike the AE model, the STL does not require a separate training before testing as it decomposes the entire data into its trend, seasonality, and residual components. In addition, the proposed STL framework is a less complex solution compared to the AE model which applies a multilayered structure for predictions. Nevertheless, the accuracy of this model is close to 99%, indicating the proposed STL solution is also robust for detecting anomalies from real sensor data in nuclear applications.

Figure 5 shows the confusion matrix AE model's testing data consisting of 2,519 values. Based on the confusion matrix, the following is a list of different performance measuring terms.

- Accuracy: Reflects the overall ability of the model to correctly determine the clean and anomaly samples.

$$Accuracy = \frac{\text{True Positive+True Negative}}{\text{Total samples}} \tag{7}$$

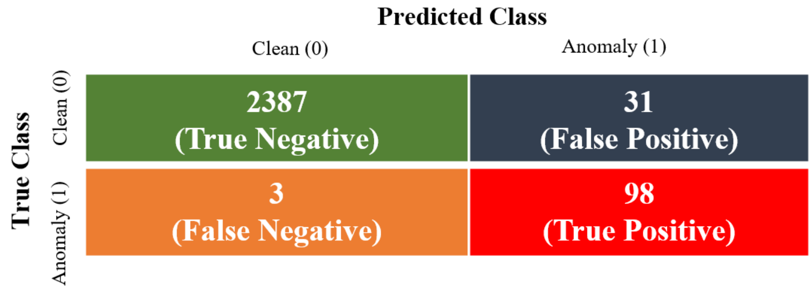

Figure 5. Confusion matrix for the AE model.

- True positive (TP): Actual anomaly samples correctly identified as anomaly. The true positive rate (TPR), which is also referred to as sensitivity or recall, is the probability that an actual positive (i.e., anomaly) is tested as positive.

$$Sensitivity\ or\ Recall = \frac{\text{TP}}{\text{TP+FN}} \tag{8}$$

- True negative (TN): Actual clean samples correctly identified as clean. True negative rate (TNR) or specificity is the probability of an actual negative (i.e., clean) is tested as negative.

$$Specificity = \frac{\text{TN}}{\text{TN+FP}} \tag{9}$$

- False positive (FP): Actual clean samples incorrectly identified as anomaly. False positive rate (FPR) is the probability of a clean sample to be incorrectly identified as an anomaly sample.

$$FPR = \frac{\text{FP}}{\text{FP+TN}} \tag{10}$$

- False negative rate (FN): The number of actual anomaly samples incorrectly identified as clean. False negative rate (FNR) refers to the probability of an anomaly sample to be incorrectly identified as clean.

$$FNR = \frac{\text{FN}}{\text{FN+TP}} \tag{11}$$

- Precision: It measures how often the model correctly predicts the anomaly class.

$$Precision = \frac{TP}{\text{FP+TP}} \tag{12}$$

- F1-score: It is a metric that combines both precision and recall. It indicates how the model works toward correctly identifying positive instances (i.e., anomalies) and minimizing false positives.

$$F1 - score = \frac{2 \text{ x Precision x Recall}}{\text{Precision+Recall}} \tag{13}$$

Precision, recall, and F1-score are the most common metrics in literature to validate the efficacy of a model. Consequently, Table 1 compares the performance of our AE and STL model in terms of these metrics.

Table I. Precision, recall, and F1-score of the AE and STL models for the imbalanced dataset.

	With Transient and Steady State			**With Steady State Only**		
Model	**Precision %**	**Recall %**	**F1-score %**	**Precision %**	**Recall %**	**F1-score %**
AE	76	97	85.2	98.9	97	97.9
STL	78.2	92	84.5	97.8	92	94.8

As shown in Table I, considering the dataset with transient and steady-state condition, the precision values of the proposed models are low even though the recall values are high. This biasing leads to an F1-score of 85.2% and 84.5%, respectively, for AE and STL models. For the AE model, the resulted biasing is not due to the imbalanced dataset, rather it is due to the false positives that were selected during the transient stage in Figure 3(b). This is confirmed by the improved precision and F1-scores observed when excluding the transient stage and focusing solely on the steady state. Since the entire dataset is imbalanced, the steady-state portion is also imbalanced. The superiority of the AE model for such an imbalanced dataset is consistent with the previous studies [10],[11] where the AE model is used as a powerful tool for imbalanced data. Note that for our considered dataset in this paper, the steady-state condition is of *greater interest* and in this region the precision is 98.9%. For the STL model, the recall value is unchanged, while precision improves due to exclusion of false alarms from the transient state (see Figure 4(b)) in the steady-state analysis.

We also examined the performance of the models on a balanced dataset by increasing the number of inserted anomalies to approximately 50%. For such a balanced dataset, the steady-state performance (which is our region of interest) is listed in Table II. Recall that the steady state data in this paper is of higher importance for informing the NPP operation decisions. For this scenario, the proposed AE shows slight improvement compared to its steady-state performance in Table I. This solidifies our finding that the biased results in Table 1 were primarily due to the false positives in the transient state, which is of less significance. On the other hand, the STL model achieves a higher precision and F1-score for a balanced dataset compared to its imbalanced counterpart. This indicates that while the neural network-based AE model can successfully detect anomalies in imbalanced datasets, the performance of the STL model can be enhanced when dealing with balanced datasets.

Table II. Precision, recall, and F1-score of the AE and STL model for the balanced dataset in steady state.

Model	**Precision %**	**Recall %**	**F1-score %**
AE	99	97.2	98
STL	99.5	99.4	99.4

In summary, the STL model is sufficient to meet the requirement for AD in balanced datasets, while the ML-based AE model offers more precise and robust detections for imbalanced datasets. Since nuclear applications often involve imbalanced data, the AE model holds an advantage in these scenarios.

4.2. Combined Performance

We further concretize the results by combining the performance of our two solutions. To do this, we assigned equal weights to our models and selected a sample as anomalous only if both solutions identified it as such. In other words, this is equivalent to using the "AND" operator on the output of the two solutions.

Figure 6(a) displays the predicted anomalies by both models for the imbalanced data. Like their individual performances, the combined performance successfully predicted 94% of the ground-truth anomaly points. In other words, the combined performance concretizes the results further and selects over 94% of the actual anomaly samples. The instances that were not selected as anomalies by one of the models were excluded from their combined selection group in Figure 6(a). Consequently, there were some instances correctly predicted as anomalous by the AE model but incorrectly identified as clean by the STL model, and vice versa. Such instances are illustrated in Figure 6(b).

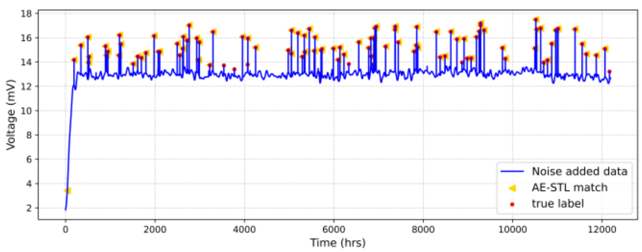

(a)

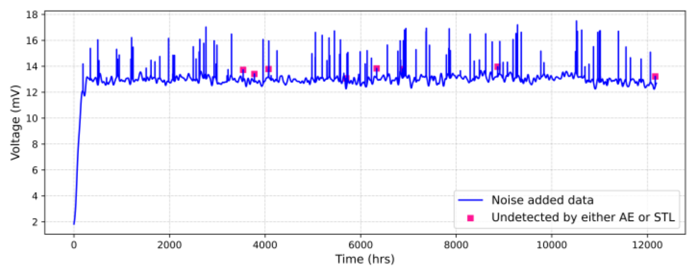

(b)

Figure 6. (a) AE-STL matched anomalies and (b) undetected anomalies by either AE or STL.

In our simulation, the combined performance missed fewer than 7% of actual anomalies. We observed that most of these missed samples were correctly identified by the AE model but ignored by the STL solution, suggesting that the AE approach is more robust. The STL model can potentially improve its performance on these samples by adjusting its parameters, such as the threshold. However, such modifications may increase the model's sensitivity to minor deviations, leading to a higher false positive rate. Despite this, the models' high accuracy in selecting true anomalies with marginal false negatives indicate their reliability for performing AD on real sensor-data-driven nuclear applications.

5. CONCLUSIONS AND FUTURE WORK

We proposed two solutions for AD from real sensor data in NPPs. Our proposed AE solution is a neural-network-based ML approach that learns through a dimensionality reduction process. Results showed that the proposed AE solution identified the actual anomalies with an accuracy of 98% for a representative imbalanced dataset. We also proposed a statistical approach—the STL model—that identified more than 92% of the actual anomalies. In addition, both models were validated on a balanced dataset for further investigation. We showed that the STL model can perform robust AD for a balanced dataset, while the AE approach outperforms for an imbalanced dataset. We further showed their combined performance by selecting only the anomaly samples when both models claimed them to be anomalous. This made the results more concrete and eliminated the reliance on a single solution. The solutions were performed assuming equal weights to both models; however, giving larger weight to the model with higher accuracy is a valid assumption and is a possible future research plan. Moreover, we intend to investigate the detected anomaly samples to find out the root cause so that they can be prevented in advance.

REFERENCES

[1] V. Chandola, A. Banerjee, and V. Kumar, “Anomaly detection: A survey,” *ACM Computing Surveys (CSUR)*, vol. 41, no. 3, pp. 1–58, July 2009, doi: 10.1145/1541880.1541882.

[2] X. Li, T. Huang, K. Cheng, Z. Qiu, and T. Sichao, “Research on anomaly detection method of nuclear power plant operation state based on unsupervised deep generative model,” *Annals of Nuclear Energy,* vol. 167, 2022, doi: 10.1016/j.anucene.2021.108785.

[3] M. Ahmed, A. N. Mahmood, and J. Hu, “A survey of network anomaly detection techniques,” *Journal of Network and Computer Applications*, vol. 60, pp. 19–31, 2016, doi: 10.1016/j.jnca.2015.11.016.

[4] I. Steinwart, D. Hush, and C. Scovel, “A Classification Framework for Anomaly Detection,” *Journal of Machine Learning Research*, vol. 6, pp. 211–232, Dec. 1, 2005.

[5] E. Eskin, A. Arnold, M. Prerau, L. Portnoy, and S. Stolfo, “A geometric framework for unsupervised anomaly detection: Detecting intrusions in unlabeled data,” *Applications of data mining in computer security*, vol. 6, pp. 77–101, 2002.

[6] M. Ester, H. P. Kriegel, J. Sander, and X. Xu, “A density-based algorithm for discovering clusters in large spatial databases with noise,” in *Proceedings of the Second International Conference on Knowledge Discovery and Data Mining (KDD’96)*, pp. 226–231, Aug. 1996.

[7] M. Ahmed and A. N. Mahmood, “A novel approach for outlier detection and clustering improvement,” In *2013 IEEE 8th Conference on Industrial Electronics and Applications (iciea)*, pp. 577-582. IEEE, 2013.

[8] J. An and S. Cho. "Variational autoencoder based anomaly detection using reconstruction probability," *Special lecture on IE,* vol 2, no. 1, pp 1-18, 2015.

[9] C. Chen, K. Yeo, B. S. Lee and C. T. Lau, “Autoencoder-based network anomaly detection,” *2018 Wireless Telecommunications Symposium (WTS)*, Phoenix, AZ, USA, 2018, pp. 1-5, doi: 10.1109/WTS.2018.8363930.

[10] V. Tomescu, G. Czibula, and Ş. Niţică. "A study on using deep autoencoders for imbalanced binary classification." *Procedia Computer Science*, pp. 119-128, vol. 192, 2021.

[11] C. Zhang, W. Gao, J. Song and J. Jiang, "An imbalanced data classification algorithm of improved autoencoder neural network," *2016 Eighth International Conference on Advanced Computational Intelligence (ICACI)*, Chiang Mai, Thailand, 2016, pp. 95-99, doi: 10.1109/ICACI.2016.7449810.

[12] M. Sakurada and T. Yairi, “Anomaly detection using autoencoders with nonlinear dimensionality reduction,’ In *Proceedings of the MLSDA 2014 2nd workshop on machine learning for sensory data analysis*, pp. 4-11. 2014.

[13] M. Ulmer, J. Zgraggen, and L. G. Huber, “A generic machine learning framework for fully-unsupervised anomaly detection with contaminated data,” *arXiv preprint arXiv:2308.13352v2,* 2023.

[14] I. C. Paschalidis and Y. Chen, “Statistical anomaly detection with sensor networks,” *ACM Transactions on Sensor Networks,* vol. 7, no. 2, pp. 1-23. 2010.

[15] N. Ye and Q. Chen, “An anomaly detection technique based on a chi-square statistic for detecting intrusions into information systems,” *Quality and reliability engineering international, vol.* 17, no. 2, pp. 105-112, 2001.

[16] C. Krügel, T. Toth, and E. Kirda, “Service specific anomaly detection for network intrusion detection,” In *Proceedings of the 2002 ACM symposium on Applied computing*, pp. 201-208. 2002.

[17] R. B. Cleveland, W. S. Cleveland, J. E. McRae, and I. Terpenning, “STL: A seasonal-trend decomposition,” *J. Off. Stat,* vol 6, no. 1, pp. 3-73, 1990.

[18] S. Lee and H. K. Kim, “Adsas: Comprehensive real-time anomaly detection system,” In *Information Security Applications: 19th International Conference*, WISA 2018, Jeju Island, Korea, August 23–25, 2018, Revised Selected Papers 19, pp. 29-41, 2019.

[19] C. Zhou and R. C. Paffenroth, “Anomaly detection with robust deep autoencoders,” In *Proceedings of the 23rd ACM SIGKDD international conference on knowledge discovery and data mining*, pp. 665-674, 2017.

[20] Y. LeCun, Y. Bengio, and G. Hinton, “Deep learning,” *Nature*, vol. 521, no. 7553, pp. 436-444, 2015.

[21] M. Thottan and C. Ji, “Anomaly detection in IP networks,” in *IEEE Transactions on Signal Processing*, vol. 51, no. 8, pp. 2191-2204, Aug. 2003, doi: 10.1109/TSP.2003.814797.

[22] X. He, Y. Li, J. Tan, B. Wu, and F. Li, “OneShotSTL: One-shot seasonal-trend decomposition for online time series anomaly detection and forecasting,” *arXiv preprint arXiv:2304.01506,* 2023.