Semana Nacional de Engenharia Nuclear e da Energia e Ciências das Radiações – VII SENCIR
Belo Horizonte, 12 a 14 de novembro de 2024

Id.: EN40

OPTIMIZING DEEP NEURAL NETWORKS FOR NUCLEAR POWER PLANT TEMPERATURE ESTIMATION IN DEEP LEARNING-BASED: A STUDY ON FEATURE IMPORTANCE AND OUTLIER DETECTION

Bernardo M. Caixeta[1], João V. S. A. Guimaraes[1], Marcelo C. Santos[1], Andressa S. Nicolau[1], Roberto Schirru[1], Da Silva M. Candeias[2], Muzitano G. Frazão[2], and Justino M. Castro[2]

bernardo@lmp.ufrj.br, sigaud@poli.ufrj.br, marcelo@lmp.ufrj.br, andressa@lmp.ufrj.br, schirru@lmp.ufrj, candeia@eletronuclear.gov.br, grazi@eletronuclear.gov.br, and justino@eletronuclear.gov.br

Key words: Nuclear Power Plant; Machine Learning; Deep Learning; Outlier Detection; Feature Importance

ABSTRACT

To support the life extension program of the Angra 1 Nuclear Power Plant (NPP), it is necessary to estimate the historical temperature exposure of equipment within the containment to assess the aging and degradation of critical components, which is essential for extending the NPP license for an additional 20 years. Since Mobile Temperature Sensors (MTSs) were installed inside the containment only in 2015, this work uses Deep Neural Networks (DNNs) to infer historical temperatures for periods before their installation. The DNNs use time series data from plant Fixed Sensors (PFS) monitored by the Integrated Computer System of Angra 1 (SICA) as input. A key step in this approach is the selection of which PFS SICA exhibits the highest correlation with the MTS temperature data for each NPP equipment. Another key aspect is handling outliers in time series data, since real-world time series, especially temperature time series data of NPPs, often contain outliers. Such outliers can distort the DNN's learning process and reduce model accuracy, potentially leading to incorrect calculations of the qualified life of the equipment. Therefore, this study conducts a comparative analysis of the impact of the DNN performance on inferring the temperature of one NPP equipment, when applying different feature importance methods, such as XGBoost, Random Forest (RF), Principal Component Analysis, and outlier detection methods, such as autoencoders, DBSCAN, and isolation forest on input time series DNN input data. The results showed that proper feature selection and outlier treatment on time series data significantly impact DNN performance. When using PCA without any outlier detection method, the DNN achieved a Mean Absolute Error (MAE) of 3.002 on the test dataset. In contrast, when Random Forest and XGBoost were used for feature selection, along with an Autoencoder for outlier detection, the DNN achieved an MAE of 0.531 on the test set. The improvement in DNN performance directly reflects the improvement in the estimation of the qualified life of NPP equipment.

1. INTRODUCTION

In the nuclear field, problems related to time series involve the continuous monitoring of critical variables such as core temperature, cooling system pressure, radiation levels, etc. This data is essential for predicting anomalies and ensuring the safe operation of nuclear facilities. Time series analysis helps detect patterns and anticipate failures in cooling and security systems, as well as monitor ambient radiation, aiding in compliance with safety standards and preventing operational risks. However, time series data often contains outliers that can affect data interpretation, potentially impacting decision-making processes and leading to inaccurate conclusions if not properly addressed.

This work presents the problem of calculating the qualified life of equipment located within the containment building of the Angra 1 Nuclear Power Plant (NPP), which was commissioned prior to 1982, and is currently undergoing a life extension process through its Environmental Qualification Program for Electrical Equipment (EQPEE) [1]. Typically, NPPs are licensed to operate for 40 years, with a possible extension of an additional 20 years. However, over time, the properties of NPP components undergo gradual changes due to aging mechanisms, which can involve one or multiple cumulative effects. These effects can potentially compromise the ability of Structures, Systems, and Components (SSCs) to fulfill their intended functions. The Angra 1 EQPEE aligns with qualification programs implemented by NPPs worldwide and adheres to the Nuclear Regulatory Commission (NRC) regulation 50.49, which pertains to the "Environmental qualification of electric equipment important to safety for nuclear power plants" [2]. The primary objective of the EQPEE is to ensure that critical equipment remains qualified to operate under both normal and adverse conditions, including the harsh environment inside the reactor containment. A key task of the EQPEE is to monitor the aging of Angra 1 SSCs to maintain NPP safety standards.

The approach implemented by the EQPEE [3], which is used in this work, utilizes temperature time series data monitored by the Integrated Computer System of Angra 1 (SICA) [4] to calculate the useful life of equipment based on the EPRI TR102167 methodology [5]. This methodology employs the Arrhenius equation to predict the rate of equipment degradation and estimate the remaining useful life, relying on historical temperature exposure data for equipment. However, acquiring this data can be particularly challenging for plants built before 1982, such as Angra 1 NPP. To overcome this challenge, the approach uses Mobile Temperature Sensors (MTS), which were installed inside the containment building in 2015 to monitor temperature conditions in areas housing Structures, Systems, and Components (SSCs). One key step in this approach is to correlate the MTS data with historical data from the Plant Fixed Sensors (PFS) monitored by SICA. The correlation process aims to identify which SICA variables exhibit the highest correlation with the MTS temperature data for each SSC within the containment building. Once these correlations are established, Deep Neural Networks (DNNs) are trained using the SICA variable data as input and the MTS temperature data as output. Once trained, the DNN can infer MTS temperatures for periods predating the installation of the MTS, effectively acting as a 'virtual' sensor for historical temperature data. This innovative approach enables comprehensive temperature monitoring, even when direct historical measurements are unavailable. In this context, a key aspect of this approach involves selecting the most relevant PFS variables (approximately 170 in total) that exhibit the strongest correlation with each MTS. Another key aspect is handling outliers in time series data, particularly in the inputs, since real-world time series often contain anomalies, especially with temperature data over extended monitoring periods. Such outliers can distort the DNN's learning process and reduce model accuracy, potentially leading to incorrect calculations of the qualified life of SSCs.

Therefore, this study proposes a comparative analysis and evaluation of the influence of feature selection and outlier treatment methods on temperature time series inference using DNNs. The analysis incorporates machine learning techniques such: XGBoost, Random Forest (RF), and Principal Component Analysis (PCA) for feature selection, as well as different outlier detection methods, such as Autoencoders (AEs), DBSCAN, Modified Tau-Thompson, and Isolation Forest. In the nuclear field, machine learning techniques like RF [6] and PCA [7] have been shown to be effective for feature selection in time series data. These methods help reduce redundant or irrelevant input variables, enabling the selection of the most significant features, which in turn enhances prediction models accuracy. Moreover, handling outliers in time series-related nuclear problems is a crucial process. So, outliers can distort the interpretation of data patterns, while improperly removing data can result in the loss of valuable information, both situations can impact prediction models performance [8].

2. OUTLIER DETECTION AND FEATURE IMPORTANCE METHODS

2.1. Outlier Detection

Outlier detection is the process of identifying data points that deviate significantly from the expected patterns or behavior within a dataset. This technique is crucial in various fields, such as cybersecurity, fraud detection, healthcare, and safety-critical systems, where recognizing anomalies can prevent errors, detect malicious activities, or uncover hidden patterns [9]. An outlier is typically defined as a data point that significantly differs from the rest of the data, either by not conforming to the expected normal behavior or by closely aligning with a known abnormal pattern [10]. Outlier detection is a complex task, despite the seemingly straightforward concept of outliers as patterns that deviate from expected normal behavior. In theory, this concept involves defining a "normal region" that encompasses all standard observations, with any data point outside this region classified as an outlier. However, defining this normal behavior or region is fraught with challenges.

One key difficulty lies in capturing the full range of possible normal behaviors within a single region, as normal behavior can vary significantly. The boundary between normal and outlier behavior is often ambiguous, leading to potential misclassifications what may be considered an outlier could be a normal instance, and vice versa. Moreover, normal behavior is dynamic and evolves over time, meaning that what is regarded as normal today may not be representative in the future. A typical outlier detection task involves identifying instances within a dataset that deviate significantly from the majority, often focusing on a fixed number of outliers. However, it is difficult to establish a universal mathematical measure for deviation that applies to all datasets and scenarios [11]. Considering the challenges inherent in outlier detection, a wide range of methods has been developed over the years to address these difficulties, spanning from traditional statistical approaches to more advanced machine learning and deep learning techniques.

2.1.1. Thompson Tau

The Thompson Tau method [12] is a classic statistical technique, proposed by R. Thompson in 1985, used to detect and exclude outliers in a dataset. It operates by calculating the mean and standard deviation of the data, and then determining a critical value (Tau) based on the number of observations and the desired significance level. Each value in the dataset is compared to the mean; if the difference between a value and the mean exceeds the product of Tau and the standard deviation, that value is classified as an outlier. If identified, the outlier is removed, and the process is repeated until no outliers remain, ensuring that atypical values do not distort the results [13].

2.1.2. Density-Based Spatial Clustering of Applications with Noise (DBSCAN)

DBSCAN[14] is a density-based clustering algorithm that excels at identifying anomalies in dense and large datasets. DBSCAN uses two user-defined parameters: the neighborhood distance (epsilon) and the minimum number of points (minpts). For a given point, points within the epsilon distance are considered neighbors, and if the number of neighbors is greater than minpts, these points form a cluster. DBSCAN classifies points as core points (with at least minpts neighbors within epsilon distance), border points (neighbors of core points but not core points themselves), or outliers (points that do not belong to any cluster). This approach enables DBSCAN to identify and classify points that do not fit into any cluster, making it useful for anomaly detection.

2.1.3. Isolation Forest

Isolation Forest[15] is an anomaly detection algorithm based on the concept of "isolation" of data. It constructs a "forest" of random decision trees, where each tree is built by randomly splitting the data space, resulting in a set of trees that represent various ways of segmenting the data. The core idea is that anomalies, being rare or different observations, are easier to isolate compared to normal data, which tends to cluster. For each tree, the algorithm randomly selects a feature and splits the data into two groups based on a random threshold for that feature. This process is repeated until each data point is isolated in a leaf of the tree or a maximum number of splits is reached. After constructing the trees, each data point's isolation is evaluated, with points that are isolated more quickly (requiring fewer splits) being more likely to be anomalies. Points that require more splits to be isolated are considered more likely to be normal. The anomaly score for each point is calculated based on the depth of the leaf where the point is isolated, with higher scores indicating anomalies [16].

2.1.4. Autoencoder

An Autoencoder (AE)[17] is a DNN model designed to learn how to reproduce its input through a process of encoding and decoding. During the encoding phase, the AE compresses the input into a lower-dimensional space, filtering out noise and outliers while retaining the most relevant features. In the decoding phase, it reconstructs the input from this compressed representation. A key feature of AEs is their ability to automatically remove outliers and noise during the encoding process, resulting in an output that closely resembles the original input but with reduced distortions. This makes AE particularly effective in filtering outliers, as demonstrated in practical applications such as correcting outliers in electrocardiogram heartbeat signals.

2.2. Feature Importance

In machine learning, feature importance involves assigning a relative score to input variables (features) based on their significance in predicting the model's output. Identifying the most influential features is crucial, especially in problems with many input variables. Moreover, accurately estimating feature importance is not only desirable but often essential before deploying models in sensitive domains [18]. However, determining feature importance is challenging due to the complex relationships between features and the target variable. Unlike simple linear models, where the influence of each feature can be directly inferred from the model coefficients, more complex models like decision trees and DNNs require specialized techniques to estimate feature importance. Therefore, common methods for assessing feature importance in complex models include evaluating changes in predictive performance when a feature is removed. If removing a feature significantly degrades the model's accuracy, that feature is deemed important. Another approach is to quantify the variance in the target variable explained by each feature, which helps in understanding the contribution of features to overall prediction variability [19]. Over the years, methods have been developed to address these challenges, particularly through advancements in machine learning techniques.

2.2.1. Random Forest

Random Forest (RF)[19], introduced by Breiman in 2001, is an ensemble learning algorithm that enhances predictive accuracy by combining multiple decision trees. It employs bagging, where decision trees are built through bootstrap sampling from the original dataset and incorporate feature randomness by selecting different subsets of features for each tree. This ensemble approach helps mitigate overfitting compared to a single decision tree, as it averages the predictions from many trees, leading to improved generalization. In terms of feature importance, RF assesses the contribution of each feature to the model's predictive performance. It evaluates the impact of each feature by measuring the change in the model’s accuracy when that feature is removed. Features that significantly reduce the model’s accuracy upon removal are deemed more important. This method provides a reliable measure of feature importance due to the aggregation of results from multiple trees within the RF model.

2.2.2. Extreme Gradient Boosting (XGBoost)

XGBoost [20] is a machine-learning algorithm, introduced by Chen and Guestrin in 2016, XGBoost builds an ensemble of decision trees using a boosting technique. Unlike bagging, which combines predictions from multiple trees built independently, boosting sequentially constructs trees, with each new tree aiming to correct the errors made by the previous ones. This iterative approach helps the model focus on difficult-to-predict instances, improving overall accuracy. For feature importance, XGBoost evaluates the contribution of each feature to the model’s predictive power by measuring how much each feature improves the model’s performance. It calculates feature importance based on metrics such as the reduction in error or gain that each feature provides when used in the trees. Features that lead to substantial improvements in model accuracy or reduction in error are considered more important.

2.2.3. Principal Component Analysis

Principal Component Analysis (PCA)[21] is a dimensionality reduction technique developed by Pearson in 1901, that transforms data into a new coordinate system where the greatest variance in the data is captured by the first few principal components. By projecting the original data onto these principal components, PCA reduces the number of dimensions while preserving as much variability as possible, which simplifies the data without losing critical information. In terms of feature importance, PCA evaluates the significance of features by examining their contribution to the principal components. Each principal component is a linear combination of the original features, and the amount of variance explained by each component reflects the importance of the features contributing to it. Features with higher loadings on principal components that explain significant variance are considered more important. PCA provides a way to identify and prioritize features based on their impact on the variance and structure of the data.

3. METHODOLOGY

The methodology for this study involved a systematic process to develop a Deep Neural Network (DNN)-based approach for inferring historical temperature data in the Angra 1 reactor building. This study aimed to compare different feature selection and outlier detection methods to determine their impact on the model's performance. The process consisted of multiple stages, including data collection, feature selection, outlier detection, data preprocessing, and model training and evaluation. The data utilized in this study was obtained from two primary sources. First, Mobile Temperature Sensors (MTS) were installed in predefined positions inside the Angra 1 reactor building in 2015 to monitor temperature conditions. Second, historical data from the Plant Fixed Sensors (PFS), monitored by the Integrated Computer System of Angra 1 (SICA), was used. The PFS data encompassed various operational variables recorded since the plant's commissioning.

To determine the most relevant input features for the DNN, the MTS data and all available SICA variables were provided to three distinct feature importance algorithms: XGBoost, Random Forest, and Principal Component Analysis (PCA). Each algorithm identified the SICA variables exhibiting the highest degree of correlation with the MTS data. From these results, the top three most important variables were selected for training the DNN, ensuring that only the most informative features were used. This comparative approach allowed the evaluation of how different feature selection methods affected the performance of the DNN model. Given the long operational timeframe and the nature of temperature data, the presence of outliers was anticipated. To mitigate the influence of outliers on model performance, four distinct outlier detection algorithms were applied during the data preprocessing stage: Thompson Tau, DBSCAN, Isolation Forest, and Autoencoder. The Autoencoder architecture used was a Convolutional Autoencoder [22] with four hidden layers: two convolutional layers with 32 and 16 filters, respectively, and two transpose convolutional layers with 16 and 32 filters, respectively. This study compared the impact of each outlier detection method on the model's performance, allowing a comprehensive evaluation of their effectiveness in enhancing data quality for training.

Following outlier detection, data gaps were addressed using a linear interpolation method, followed by forward and backward filling to ensure no missing values remained in the dataset. This preprocessing step was crucial for ensuring data quality and continuity, which are vital for effective DNN training. The DNN architecture consisted of five hidden layers with 100 neurons per layer. The model was trained for 500 epochs to ensure convergence. To maintain consistency and reproducibility, all DNNs were trained using the same dataset split: 50% for training, 20% for validation, and 30% for testing. The Mean Absolute Error (MAE) metric was used to evaluate model performance, measuring the error between the actual and predicted temperature values. By comparing the MAE across different feature selection and outlier detection methods, the study assessed their relative contributions to model accuracy.

4. RESULTS

After selecting and processing the data, the Deep Neural Networks (DNNs) were trained, and the results are summarized in Table 1. Table 1 presents the MAE obtained by the DNNs, considering all combinations of variable selection methods (three methods) and outlier detection methods (four methods), resulting in a total of 12 distinct DNN models. The evaluation aimed to determine the influence of variable selection and outlier detection on the DNN's ability to accurately infer historical temperature data, which is critical for assessing the useful life of NPP components.

The analysis of input variable selection methods showed that both XGBoost and Random Forest identified the same set of SICA PFS variables (TI5700, TI5701, TI5702) as the optimal input for the DNN. Using these variables without applying any outlier detection resulted in a MAE of 1.033 on the test set. In contrast, the PCA method selected a different set of input variables (FT474D, TE05, TE23), and when used without any outlier detection, the resulting DNN had a significantly higher MAE of 3.002, approximately three times greater. This demonstrates the substantial impact that choosing an appropriate variable selection method can have on the DNN's performance.

Furthermore, the effect of outlier detection methods on the best-performing set of input variables (TI5700, TI5701, TI5702) was analyzed. Most outlier detection methods improved the DNN's performance. The DNN trained using input data processed with the Isolation Forest method achieved a MAE of 1.031, while the DBSCAN method resulted in a MAE of 0.729. The Modified Tau-Thompson and Autoencoder (AE) methods achieved the lowest MAE values, 0.659 and 0.531, respectively. These findings emphasize the significant benefits of employing an appropriate outlier detection method to preprocess input data, resulting in more accurate DNN predictions.

The improvement in DNN performance directly reflects on the estimation of the qualified life of NPP equipment. The trained DNNs are used to infer the historical temperature exposure of equipment throughout its operation in the NPP, and the inferred temperature values are then used in the Arrhenius equation to predict the rate of material degradation and estimate the remaining useful life of the equipment. This estimation is essential for assessing whether the equipment can continue to safely operate within the NPP and for determining how many more years of service life can be expected before replacement is necessary. Therefore, the greater the precision of the DNN in estimating temperature, the more accurate the calculation of the equipment's qualified life. This precise calculation increases the reliability of the approach, which is a key requirement in the licensing process for extending the operational life of the nuclear plant.

Table 1: Results obtained for the DNNs trained for inference of MTS on the Position 1 24

Selection Method	Outlier Method	MTS Position	SICA PFS Variables	Training Error	Validation Error	Test Error
XGBoost / Random Forest	-	1 24	TI5700, TI5701, TI5702	1.251	0.245	1.033
XGBoost / Random Forest	TAU	1 24	TI5700, TI5701, TI5702	1.498	0.232	0.659
XGBoost / Random Forest	AE	1 24	TI5700, TI5701, TI5703	0.410	0.227	0.531
XGBoost / Random Forest	DBSCAN	1 24	TI5700, TI5701, TI5703	0.496	0.216	0.729
XGBoost / Random Forest	Isolation Forest	1 24	TI5700, TI5701, TI5703	1.31	0.224	1.031
PCA	-	1 24	FT474D, TE05, TE23	1.596	1.265	3.012

PCA	TAU	1 24	FT474D, TE05, TE23	1.519	1.249	2.302
PCA	AE	1 24	FT474D, TE05, TE23	1.614	1.242	2.375
PCA	DBSCAN	1 24	FT474D, TE05, TE23	1.491	1.020	2.740
PCA	Isolation Forest	1 24	FT474D, TE05, TE23	2.469	1.043	2.843

5. CONCLUSION

This study conducted a comparative analysis of the impact of feature selection and outlier treatment methods on the accuracy of inferred temperature time series for NPP equipment within the containment of Angra 1. In this work approach DNNs were employed as virtual sensors of pasta, enabling the inference of historical temperature profile exposure of NPP equipment at critical locations inside the plant containment. By estimating the temperatures experienced by key Structures, Systems, and Components (SSCs), the DNNs allow the assessment of the remaining useful life of the equipment, this assessment is an essential part of the processing of evaluating the feasibility of extending the operational lifespan of the Angra 1 NPP.

The results highlighted that the choice of feature selection and outlier detection techniques for the input time series data has a significant influence on DNN predictive performance. Improving temperature prediction consequently enhances the accuracy of calculating the qualified life of the equipment, supporting the plant life extension program. Among the variable selection methods evaluated, Principal Component Analysis (PCA) yielded suboptimal results, with a Mean Absolute Error (MAE) of 3.002, whereas both Random Forest and XGBoost identified a more effective set of input variables, resulting in an MAE of 1.033. Furthermore, the application of outlier detection methods substantially enhanced the predictive accuracy of the DNN in this set of variables, with the Autoencoder-based approach achieving the lowest MAE of 0.531. These findings underscore the critical importance of appropriate input variable selection and robust outlier handling in optimizing DNN performance.

The enhanced accuracy of DNN-based temperature inference has direct implications for estimating the qualified life of Angra 1 SSCs, which is a crucial aspect of the Angra 1 Environmental Qualification Program for Electrical Equipment (EQPEE). Accurate historical temperature estimation is essential for predicting material degradation rates using the Arrhenius equation, thereby enabling precise calculations of the remaining service life of equipment in the containment of the NPP. Such precision enhances the reliability of life extension evaluations, which is a fundamental requirement in the licensing process for extending the operational life of Angra 1 by an additional 20 years.

The findings of this study reveal the importance of integrating machine learning methodologies in the nuclear area, particularly for the life extension of aging nuclear power plants. The demonstrated efficacy of the Random Forest/XGBoost and Autoencoder-based approaches suggests promising avenues for future research, such as exploring more sophisticated deep learning techniques for feature selection and outlier detection, including Variational Autoencoders (VAE) and Generative Adversarial Networks (GAN). Applying these advanced methods could further enhance the accuracy of DNN-based temperature inference, ultimately contributing to safer NPP equipment operations during NPP extended service life.

ACKNOWLEDGMENTS

The authors would like to acknowledge: CAPES (Coordination for the Improvement of Higher Education Personnel), CNPq (National Council for Scientific and Technological Development), FAPERJ (Carlos Chagas Filho Foundation for Research Support of the State of Rio de Janeiro), INCT (National Institutes of Science and Technology of Innovative Reactors) and Eletronuclear for financial support.

REFERENCES

[1] Muzitano et al. 2017. Environmental qualification program of electric equipment for Angra 1. In Proceedings of the 2017 International Nuclear Atlantic Conference (INAC 2017), Belo Horizonte, MG, Brazil. Associação Brasileira de Energia Nuclear (ABEN).

[2] NRC, 10 CFR 50.49 Environmental qualification of electric equipment important to safety for nuclear power plants, United States Nuclear Regulatory Commission, USA (updated 2015).

[3] Nicolau et al. 2023. Deep neural networks for estimation of temperature values for thermal ageing evaluation of nuclear power plant equipment. Progress in Nuclear Energy. 153, 104542.

[4] Schirru, R., Pereira C.M.N.A. (2004). A Real –Time Artificially Intelligent Monitoring System for Nuclear Power Plants Operators Support. Real-Time Systems, 27, 71-83.

[5] EPRI TR 1021067, 2010. Plant Support Engineering: Nuclear Power Plant Equipment Qualification Reference Manual, Revision 1 – Final Report. September.

[6] Desterro, F. S. M., Pinheiro, V. H. C., Pereira, C. M. N. A., & Schirru, R. (2023). Prediction of LOCA's break size and location based on random forest and multi-tasking deep neural network. Nuclear Engineering and Design, 415, 112711. https://doi.org/10.1016/j.nucengdes.2023.112711.

[7] De Souza et al. (2024). Accident classification methodology with don't know response for PWR nuclear reactors using the cuckoo optimization algorithm and principal component analysis method. Nuclear Engineering and Design, 423, 113200.

[8] Valdetaro, Eduardo Damianik, and Schirru, Roberto. "Simultaneous Model Selection, Robust Data Reconciliation and Outlier Detection with Swarm Intelligence in a Thermal Reactor Power Calculation." Annals of Nuclear Energy, vol. 38, pp. 1820-1832, 2011.

[9] Singh, K., & Upadhyaya, S. (2012). Outlier detection: applications and techniques. International Journal of Computer Science Issues (IJCSI), 9(1), 307.

[10] Varun Chandola et al. 2009. Anomaly detection: A survey. ACM Comput. Surv. 41, 3 (2009), 15.

[11] Boukerche, A.; Zheng, L.; Alfandi, O. Outlier Detection: Methods, Models, and Classification. ACM Comput. Surv. 2021, 53, 55.

[12] Thompson, R. "A Note on Restricted Maximum Likelihood Estimation with an Alternative Outlier Model." Journal of the Royal Statistical Society. Series B (Methodological), vol. 47, no. 1, 1985, pp. 53–55. JSTOR

[13] Domański, P. D., (2020). "Statistical outlier labelling – a comparative study," 2020 7th International Conference on Control, Decision and Information Technologies (CoDIT), Prague, Czech Republic, pp. 439-444.

[14] M. Çelik et al, "Anomaly detection in temperature data using DBSCAN algorithm," 2011 International Symposium on Innovations in Intelligent Systems and Applications, Istanbul, Turkey, 2011, pp. 91-95.

[15] M. Ester, H.-P. Kriegel, J. Sander, and X. Xu, "A density based algorithm for discovering clusters in large spatial databases with noise," in KDD-96 Proceedings, pp. 226-231, 1996.

[16] Fei Tony Liu, Kai Ming Ting, and Zhi-Hua Zhou. 2008. Isolation forest. In 2008 eighth ieee international conference on data mining. IEEE, 413–422.

[17] Karpinski, M., et al. Autoencoder Neural Networks for Outlier Correction in ECG- Based Biometric Identification, 2018 IEEE 4th International Symposium on Wireless Systems within the International Conferences on Intelligent Data Acquisition and Advanced Computing Systems (IDAACS-SWS), Lviv, 2018, pp. 210-215.

[18] S. Hooker, D. Erhan, P.-J Kindermans and B. Kim, "A benchmark for interpretability methods in deep neural networks", Proc. Adv. Neural Inf. Process. Syst., 2019.

[19] Breiman, L. (2001). Random forests. Machine Learning, 45(1), 5-32.

[20] Chen, T., & Guestrin, C. (2016). XGBoost: A scalable tree boosting system. In Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (pp. 785-794).

[21] Pearson, K. (1901) On Lines and Planes of Closest Fit to Systems of Points in Space. The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science, 2, 559-572.

[22] Thill, M., Konen, W., Wang, H., & Bäck, T. (2021). Temporal convolutional autoencoder for unsupervised anomaly detection in time series. Applied Soft Computing, 112, 107751.