SYSTEMATIC REVIEW OF MODELS USED TO HANDLE CLASS IMBALANCE IN ANOMALY DETECTION FOR ENERGY CONSUMPTION

David Kaimenyi Marangu[1], Stephen Thiiru Njenga[2] and Rachael Njeri Ndung'u[1]

[1]Department of Information Technology, Murang'a University of Technology, Kenya
[2]Department of Computer Science, Murang'a University of Technology, Kenya

ABSTRACT

The widespread integration of Smart technologies into energy consumption systems has brought about a transformative shift in monitoring and managing electricity usage. The imbalanced nature of anomaly data often results in suboptimal performance in detecting rare anomalies. This literature review analyzes models designed to address this challenge. The methodology involves a systematic literature review based on the five-step framework proposed by Khan, encompassing framing research questions, identifying relevant literature, assessing article quality, conducting a critical review, and interpreting results. The findings show that classical machine learning models like Support Vector Machines (SVM) and Random Forests (RF) are commonly used. In conclusion, classical machine learning models like SVM and RF struggle to recognize rare anomalies, while deep learning models, notably Convolutional Neural Networks (CNNs) and Long Short-Term Memory (LSTM), show promise for automatically learning elaborate representations and improving performance while dealing with class imbalance.

KEYWORDS

Anomaly Detection, Class Imbalance, Energy Consumption, models

1. INTRODUCTION

The widespread integration of smart technologies into energy consumption systems has brought about a transformative shift in the monitoring and management of electricity usage [1]. With the increasing deployment of smart grids and advanced metering infrastructure, a substantial amount of data is generated, offering a unique opportunity for enhancing anomaly detection in energy consumption patterns [2]. However, this advantage is accompanied by challenges inherent in the highly imbalanced nature of anomaly data, where instances of abnormal behavior are outnumbered by their normal counterparts [3]. The class imbalance issue poses a significant obstacle in developing effective anomaly detection models [3]. Traditional machine learning algorithms, when confronted with imbalanced datasets, tend to favour the majority class, resulting in suboptimal performance in identifying rare anomalies crucial for maintaining the stability and efficiency of energy systems [1]. In response to this challenge, recent research has increasingly focused on addressing class imbalance, proposing a variety of models and techniques to enhance the robustness of anomaly detection in energy consumption data.

This literature review aims to comprehensively examine models specifically tailored to handle the class imbalance challenge in the context of anomaly detection for energy consumption. By synthesizing and critically analyzing the existing body of knowledge, this review seeks to elucidate the strengths, limitations, and potential avenues for improvement in the current landscape of anomaly detection models applied to imbalanced energy consumption datasets. The

DOI:10.5121/ijaia.2024.15304

review focuses on a spectrum of models, scrutinizing each model's efficacy in addressing class imbalance concerning energy consumption anomaly detection. Factors such as algorithmic adaptability, scalability, and suitability for real-world deployment are considered. Moreover, the review explores how various research endeavors have leveraged transfer learning, ensemble methods, and novel loss functions to mitigate the impact of class imbalance on the accuracy of anomaly detection. The importance of oversampling techniques, especially in deep learning models, is emphasized, along with the ongoing challenges in capturing sparse anomalies and learning from multimodal distributions. The exploration of ensemble methods, transfer learning, and specialized architectures showcases continuous evolution in addressing class imbalance challenges.

2. METHODOLOGY

The examination of models addressing the class imbalance challenge in anomaly detection for energy consumption involves a critical analysis aimed at enhancing the accuracy and efficacy of anomaly detection systems in dynamic energy landscapes. With the growing complexity and volume of energy consumption data, the inherent class imbalance between normal and anomalous instances poses a significant hurdle for models. This review explores a range of innovative approaches, spanning classical techniques to cutting-edge deep learning architectures, with the goal of unravelling strengths, limitations, and advancements.

Through a comprehensive analysis of existing literature, this review aims to provide insights informing the development of state-of-the-art anomaly detection frameworks tailored to the unique challenges posed by imbalanced energy datasets. The review of models addressing the class imbalance challenge in anomaly detection for energy consumption follows Khan, Kunz, Kleijnen, and Antes' five-step process for performing a critical review [4]. This methodology includes framing the review questions, identifying relevant literature, assessing article quality, conducting a critical review summary, and interpreting results. The research questions are first explained, followed by a search of relevant literature. The criteria for quality assessment are then outlined. The five-step methodology concludes with a critical review summary and a discussion of the results. Following the review process, proposals for future study directions are provided in the last section, giving guidance and foundation for further study on addressing the class imbalance challenge in anomaly detection for energy consumption. The research question for this study is formulated as follows:

RQ1: What models are used to address the class imbalance challenge in anomaly detection in energy consumption?

In pursuit of this objective, a concise systematic literature review and analysis are carried out utilizing a desktop research design. Beyond pinpointing gaps within existing models, the systematic review process of the literature also plays a crucial role in crafting a conceptual framework for this study. These findings further inform the experimental design methodology employed in this research. The strategy employed for identifying pertinent literature involves a meta-search approach, concentrating on models addressing the class imbalance challenge in anomaly detection for energy consumption. Key search terms include "class imbalance," "anomaly detection," and "energy consumption." Thorough searches are performed across a variety of databases, including Google Scholar, ACM, Research Gate, IEEE Xplore digital library, and Springer, among others. A total of 27 articles, published between 2016 and 2023, are successfully identified through this comprehensive search [4].

The quality assessment and selection criteria place a special emphasis on examining the abstracts of the 27 identified papers. Following the assessment, twelve papers are deemed relevant to the models designed for addressing the class imbalance challenge in anomaly detection for energy consumption. These selected papers undergo thorough review, analysis, and discussion, and their outcomes are presented in the results section [4].

The inclusion criteria applied to gather pertinent papers for this study encompass articles addressing the class imbalance challenge in anomaly detection for energy consumption. Specifically, selected papers are required to describe original work, featuring developed and validated models designed to tackle the class imbalance issue in anomaly detection for energy consumption. The chosen papers included in this research review article are limited to those published within the last five years, ensuring the most recent insights into the current state of research on the class imbalance problem. This timeframe is deemed crucial for informing recommendations and guiding future research endeavours in this area [4]. Exclusion criteria were guided by the research question of the study and the key search terms

3. RESULTS AND DISCUSSION

The results section elucidates the diverse models employed to confront the class imbalance challenge in anomaly detection for energy consumption. A comprehensive overview of the models, their training methods, and succinct descriptions, encompassing strengths and limitations, is provided. Additionally, an overview of the algorithmic performance and statistical measures, along with specific scores for each model, is presented. These findings are intricately aligned with the research question posed in this study, offering a comprehensive understanding of the models' effectiveness and performance in addressing class imbalance challenges.

According to [5] a combined genetic optimization with the AdaBoost ensemble model was proposed for detecting anomalies in buildings' electricity consumption. This is an ensemble model for electricity theft detection based on genetic optimization. First, the synthetic minority oversampling technique (SMOTE) training approach is used to create synthetic samples with a distribution that is roughly similar to that of real samples of electricity theft. Then, Principal Component Analysis (PCA) is used to extract aspects of anomalous electricity consumption with the help of dimension reduction. Ultimately, “an AdaBoost-based ensemble deep learning network is built to extract implicit information from continuous time series data and identify unusual electricity usage among labelled users. Additionally, the deep neural network's hyperparameters are optimized using a genetic algorithm (GA). The results demonstrate that this model is superior to other models, such as Support Vector Machine (SVM), Random Forest (RF), and the traditional artificial neural network (ANN), in terms of sensitivity and area under the curve (AUC), using data about user electricity consumption collected from smart meters owned by the State Grid Corporation of China” [5]. However, this approach still relies, in part, on a misjudgment of the relatively small percentage of users with less-than-obvious traits. For instance, in instances when only a small amount of electricity was involved, the theft appeared to be very effectively hidden, and load fluctuations brought on by business trips may enable regular users to be mistaken for those who are committing theft. To address this, a combination of improved hyperparameter optimization techniques, training procedures, and feature extraction technologies may be used.

According to [6] “an adaptive synthesis to handle imbalanced big data with deep Siamese network for electricity theft detection in smart grids was proposed”. The model utilizes adaptive synthesis (ADASYN) to address the issue of data-imbalanced classes. Deep Siamese networks (DSNs) that combine convolutional neural networks (CNN) and long-short-term memory (LSTM) may distinguish between the features of truthful and dishonest customers. The CNN module is tasked with extracting features from weekly energy use profiles, whilst the LSTM module handles sequence learning. Lastly, the DSN makes a judgement after evaluating the

CNN-LSTM's shared features. Data from real-time smart meters is subjected to various train-test ratios for data analytics. The simulation results show that the suggested model is effective in terms of area under the curve, F1-Score, precision, and recall. However, Deep Siamese Networks focus on learning local patterns within pairs of instances and may face challenges in capturing global patterns across multiple features in energy consumption data. Anomalies might exhibit complex interactions that are not easily captured by Siamese Networks. Moreover, CNNs are primarily designed for spatial data, and they might not naturally capture temporal dynamics in time series data, which is common in energy consumption analysis. Anomalies often exhibit temporal patterns that may not be fully captured by standard CNN architectures. Since LSTMs are meant to capture sequential dependencies, they could have trouble capturing a global pattern across multiple features in energy consumption data. Anomalies might exhibit complex interactions that are not easily captured by LSTMs.

“Anomaly Prediction in Electricity Consumption Using a Combination of Machine Learning Techniques” was proposed by [7]. In this model, the Isolation Forest algorithm is used to classify smart meter readings of power usage as normal or abnormal. It produces a series of data with variable lengths. Random Forest and Decision Tree algorithms are used to predict the occurrence of power abnormalities based on data sequence. The experiment results indicate that the model accurately forecasts aberrant status 30 minutes ahead. There is no significant difference in performance between Random Forest and Decision Tree for varied smart meter readings, dataset sizes, and data sequence lengths. This method depicts an alternate approach that is capable of auto labelling normal and abnormal data, dealing with the sequence of label data in the prediction process while avoiding the dynamic behavior of the power consumption data. However, Random Forests can be biased toward the majority class when the dataset is imbalanced. The majority class may dominate the learning process, leading to a model that performs well in normal instances but struggles to identify anomalies effectively. During training, Random Forests' default behaviour may cause them to become less sensitive to the minority class (anomalies)The algorithm tends to prioritize accuracy, and if anomalies are rare, the model may not adequately capture their patterns. Random Forests may face difficulties in learning rare and complex anomaly patterns, especially when anomalies manifest in diverse ways. The ensemble may not sufficiently focus on capturing the nuances of rare instances.

Electrical Network Fraud Detection Using a CNN and LSTM Model was proposed by [8]. This is a Deep Learning model that combines a Convolutional Neural Network Distributed in Time (CNN) with Long Short-Term Memory (LSTM). The purpose is to better understand and assess current fraud detection techniques, as well as to examine the applicability and efficiency of CNN + LSTM models in fraud detection. The model design starts with a time-distributed wrapper layer that encapsulates a 1D convolutional layer with 64 filters and 32 kernels of size 3, followed by a max pooling operation. The model then flattened these features and fed 50 units to the LSTM layer. The output layer was made up of a single unit that used a sigmoid activation function to predict the binary outcome. The Adam optimizer and binary cross-entropy loss function were used to construct the model at a learning rate of 0.01. The model was trained using a reshaped version of the balanced training dataset, with each training epoch comprising a 20% validation set. Dropout was initially thought of as a regularization strategy, however it was discovered that the model performed better without it. To balance the data, oversampling on the irregular ones was done. However, In datasets with a limited number of anomaly instances, CNNs may not have sufficient exposure to these rare instances during training. As a result, the model may struggle to accurately learn the features associated with anomalies. Moreover, CNNs may become biased toward the majority class in imbalanced datasets [8]. The model might achieve high accuracy by predicting the majority class, while neglecting the minority class (anomalies), leading to reduced effectiveness in anomaly detection. However, Oversampling can lead to overfitting, especially when the same minority class instances are repeatedly sampled. The model may memorize the oversampled instances rather than learning generalizable patterns, resulting in reduced performance on new, unseen data. Moreover, Oversampling increases the representation of minority class instances, but it may come at the cost of losing diversity within the minority class. If anomalies exhibit various patterns, oversampling methods that replicate specific instances may not capture the full range of anomaly characteristics.

Electricity Theft Detection Using Supervised Learning Techniques on Smart Meter Data was proposed by [9]. The electricity data in this model are pre-processed using interpolation, the three sigma rule, and normalization methods. The Adasyn algorithm is used to address the issue of class imbalance. First and foremost, it increases the number of minority class samples in the dataset.

The balanced data is fed into a Visual Geometry Group (VGG-16) module, which identifies unusual patterns in electricity consumption. For classification, the Firefly Algorithm employs the Extreme Gradient Boosting (FA-XGBoost) technique. Simulations are used to carry out performance evaluation. Comparative analysis is done using other models, such as Support Vector Machine (SVM), Convolution Neural Network (CNN), and Logistic Regression (LR). The metrics used include F1-score, Matthews Correlation Coefficient (MCC), Receiving Operating Characteristics Area Under Curve (ROC-AUC), Precision Recall Area Under Curve (PR-AUC), and precision. First, the simulation results show that the FA-XGboost classifier performed better when utilizing the proposed Adasyn technique, with an F1-score of 93.7%, precision of 92.6%, and recall of 97%. Second, the VGG-16 module achieved accuracy of 87.2% and 83.5% on training and testing data, respectively, resulting in improved overall performance. Third, the suggested FA-XGBoost accurately identifies actual electrical thieves, with a recall of 97%. Furthermore, the model outperforms other state-of-the-art models in terms of handling big time series data and accurate classification. However, XGBoost based models may exhibit a bias toward the majority class in imbalanced datasets. The algorithm is optimized for overall accuracy, and in the presence of imbalanced classes, it may prioritize learning patterns from the majority class while neglecting the minority class (anomalies). Moreover, XGBoost's default behavior may result in limited sensitivity to the minority class. The model may not effectively capture the nuances of anomalies, especially when they are rare and have diverse patterns. In such cases, effective assignment of weights to different classes during the training process of the XGBoost models could definitely lead to optimal anomaly detection performance.

Electricity Theft Detection in Smart Grids Based on Deep Neural Network was proposed by [10]. Data interpolation and synthetic data generation processes address class imbalance challenges. Features from both time and frequency domains are analyzed and compared. Principal Component Analysis is used for feature extractions. A Bayesian and adaptive moment estimation optimizers are used for hyperparameter optimization. Validation produced a 97% area under the curve (AUC), which is 1% higher than the best AUC in existing works, as well as 91.8% accuracy, which is second highest on the benchmark. However, when using PCA in this method, the PCA aims to capture the maximum variance in the data, which may not align with the discriminative patterns associated with anomalies. In imbalanced datasets, anomalies might be overshadowed by the dominant normal patterns, limiting the effectiveness of PCA in anomaly detection. Moreover, Deep neural networks are susceptible to overfitting, especially in the presence of imbalanced classes. The model may memorize noise or specific patterns in the training data, leading to reduced generalization on new, unseen instances.

An Ensemble Machine Learning model for detecting energy theft was proposed by [11]. This paper proposes ensemble machine learning (ML) models based on customer consumption patterns to identify energy theft in smart grids. Ensemble Machine Learning models are meta-algorithms that blend several ML approaches intelligently into a single predictive model to reduce variation and bias. The false positive and detection rates of the adaptive boosting, category boosting, extreme-boosting, light boosting, random forest, and extra trees algorithms were investigated. To address overfitting, the statistical technique of minority oversampling was utilized. Results show that bagging models outperform alternative methods. The random forest and extra trees models obtained the greatest AUC score of 0.90. The proposed bagging methods outperform the others according to the precision analysis. However, effective hyper-parameter tuning to determine the optimum ensemble weights, although complex, is critical to improve anomaly detection performance and classification accuracy of the ensemble ML model. Moreover, ensemble models can be sensitive to noisy or irrelevant features in the training data. If anomalies introduce noise, the ensemble may incorporate this noise and lead to suboptimal performance of the model. The effectiveness of an ensemble relies on the diversity among its base models. If the base models are too similar, the ensemble may not capture a wide range of anomaly patterns, limiting its ability to address class imbalance effectively

"Fault detection based on one-class deep learning for manufacturing applications limited to an imbalanced database" proposed by [12]. This model consists of three sub-modules: residual computation, one-class classification with isolation forest and one-class support vector machine, and deep learning-based time-series prediction. For time-series prediction, four distinct networks were trained using just the production success cases: Multilayer Perception (MLP), Residual Networks (ResNet), Long Short-Term Memory (LSTM), and ResNet-LSTM. The one-class classification is constructed using the residuals of the deep learning prediction as an elaborated feature. The fault detection module is evaluated using actual mass production data from a diecasting process. The one-class classification's overall accuracy increased greatly from 90.0% to 96.0%. The accuracy of detecting manufacturing defects increased from 84.0% to 96.0%. The area under the receiver operating characteristics (AUROC) also rose from 87.56% to 98.96%. While ResNet–LSTM yielded better results for guaranteeing the success of production, ResNet had the best performance for detecting production failures. One-class deep learning is a potential method for extracting key features from time series data to build a one-class defect detection module.

However, the use of the one class support vector machine in this model poses a number of challenges. One class support vector machine has hyperparameters that need to be tuned, such as the kernel parameters and the "nu" parameter, which controls the upper bound on the fraction of margin errors and support vectors. Finding optimal hyperparameters may be challenging and can impact the model's performance, especially where there is a class imbalance. One class support vector machine (OCSVM) assumes an unimodal distribution of normal instances. If the distribution of normal data in energy consumption is multimodal, OCSVM may struggle to capture the complexity of the data and may not effectively model anomalies that deviate from these modes. Moreover, the performance of OCSVM can depend heavily on the choice of features and the quality of feature engineering. If the features do not adequately represent the characteristics of anomalies, the model's performance may be suboptimal.

"Imbalanced learning algorithm based intelligent abnormal electricity consumption detection" was proposed by [13]. Using K-means clustering and the Synthetic Minority Oversampling Technique (SMOTE) in conjunction with an artificial neural network (ANN) trained by a kernel extreme learning machine (KELM), this method first examines two efficient abnormal electricity consumption detection algorithms from the perspectives of data balancing and data weighting, respectively. Next, an improved deep representation network-based ELM (EH-DrELM) and an improved multiclass AdaBoost imbalanced learning algorithm (AdaBoost-ID) are examined, along with the deep weighted ELM (DWELM). However, the use of the SMOTE in this method poses several challenges. Features are assumed to be independent of one another by SMOTE. In scenarios where features are correlated, the synthetic instances generated by SMOTE may not accurately represent the underlying distribution of the minority class. SMOTE may struggle to effectively capture rare and complex patterns associated with anomalies, especially if anomalies are sporadic or exhibit diverse characteristics. The synthetic instances may not adequately represent the true complexity of anomalies. Moreover, SMOTE can be sensitive to noisy data. If the minority class contains noisy instances or outliers that do not represent true anomalies, the synthetic instances generated by SMOTE may incorporate this noise, leading to suboptimal performance.

Intelligent Anomaly Detection Method of Gateway Electrical Energy Metering Devices using Deep Learning was proposed by [14]. This is a hybrid deep-learning model that uses stacked autoencoders (SAE) and long short-term memory (LSTM) to automatically detect anomalous occurrences in gateway electrical energy metering devices. The proposed model, known as the SAE-LSTM model, first employs SAE to extract deep latent features from three phase voltage data collected from the gateway electrical energy metering devices, and then uses LSTM to separate anomalous occurrences based on the extracted deep latent features. The SAE-LSTM model can effectively highlight temporal information in electrical data, increasing the accuracy of anomaly detection. The simulation demonstrates the SAE-LSTM model's advantages in anomaly detection at various signal-to-noise ratios. The experimental results of real datasets suggest that it is appropriate for anomaly detection of gateway electrical metering devices in practical scenarios. However, Stacked Autoencoders (SAEs) rely on the ability to represent anomalies effectively in the learned latent space. If anomalies in energy consumption data do not have clear and distinguishable patterns, or if they are highly variable, SAEs may struggle to learn a meaningful representation. SAEs provide a static representation based on the training data. If anomalies exhibit dynamic or evolving patterns over time, the model may not adapt well to changes, potentially missing novel anomaly patterns. Additionally, in the latent space, stacked autoencoders assume a unimodal distribution. Multimodal or complexly structured distributions of normal and anomalous instances may make it difficult for SAEs to fully capture the diversity of the data.

[15] presented real-world anomaly detection using digital twin systems while being weakly supervised. This model uses a Digital Twin to generate a training dataset, which simulates the normal operation of the machinery, along with a small set of labelled anomalous measurements from the real machinery. Specifically, a neural architecture based on the Siamese Autoencoders (SAE) and a clustering-based strategy dubbed cluster centres (CC) are presented. These approaches are designed for poorly supervised settings with a very small labelled data sample. Using an application to a real-world dataset from a facility monitoring system, the performance of the proposed model is compared using many performance metrics against different state-of-the-art anomaly detection algorithms. Also, the influence of hyperparameters related to feature extraction and network architecture is examined. The results show that the proposed SAE-based solutions perform very substantially better than the most advanced anomaly detection techniques over a wide range of hyperparameter settings on all performance metrics. However, Siamese Autoencoders may require a sufficient number of anomaly examples during training to learn a robust similarity metric. In scenarios with a scarcity of anomaly instances, the model's ability to generalize to new anomaly patterns may be limited. In datasets with extreme class imbalance where anomalies are significantly outnumbered by normal instances, Siamese Autoencoders may prioritize learning the dominant normal patterns and struggle to effectively represent the minority class. Moreover, Siamese Autoencoders may struggle to effectively capture sparse anomalies that are scattered across the feature space. The learned similarity metric may not highlight the specific features associated with sparse anomalies.

Self-Adaption AAE-GAN for Aluminum Electrolytic Cell Anomaly Detection was proposed by [16]. In this model, the ability of Generative Adversarial Network (GAN) is utilized to model complicated high-dimensional image distribution. This model is based on adaptive changes of input samples and is known as the self-adaption AAE-GAN network. This method tackles the aforementioned issues by transforming multi-dimensional time series data into a two-dimensional matrix and using just normal samples for training. The method makes use of encoders and decoders to create discriminators and generators. In order to completely reflect the encoder's mapping capacity, the generator and discriminator go through a contentious and cooperative training process. The size of the reconstruction difference is used to assess if the sample is abnormal during the anomaly detection stage. Results from the experiment indicate that this model has extremely high detection accuracy and speed. GANs are susceptible to mode collapse, which occurs when the generator generates a small variety of samples, which leaves the anomaly space underrepresented. This can be problematic for detecting diverse anomaly patterns in energy consumption data. GANs may struggle to control the generation of anomalies, especially when anomalies are rare and sporadic. Ensuring that the generator focuses on generating realistic and meaningful anomalies can be challenging. Moreover, the training dynamics of GANs may be affected by class imbalance, especially when anomalies area minority class. The generator might prioritize generating majority class samples, leading to an imbalance in the generated data.

The literature review findings underscore the prevalence of Support Vector Machines and Random Forests as common classical machine learning models for addressing class imbalance challenges in anomaly detection within energy consumption data. In contrast, the growing popularity of deep learning models, particularly Convolutional Neural Networks (CNNs) and Long Short-Term Memory (LSTM), demonstrates their effectiveness. The prevailing training method involves oversampling the minority class, while the evaluation metrics encompass accuracy, Area under the curve (AUC), F1 score, and Precision & recall.

Despite the success of classical machine-learning models, deep learning models demonstrate superior performance, attributing their success to their inherent capacity to automatically learn intricate representations, addressing class imbalance challenges. However, challenges persist. Support Vector Machines struggle with sparse anomalies due to their focus on maximizing the margin between classes, while Random Forests may falter in detecting sporadic or infrequent anomalies. Deep learning models, although potent, face hurdles in capturing sparse anomalies, particularly with Convolutional Neural Networks, and struggle with learning from multimodal distributions.

Moving forward, tuning deep neural networks to capture complex patterns in imbalanced data offers a promising avenue. Additionally, adopting effective transfer learning strategies and ensemble methods holds potential. Transfer learning leverages knowledge from a source task to improve anomaly detection in imbalanced scenarios, while ensembles mitigate overfitting by incorporating diverse model perspectives.

The analysis of evaluation metrics highlights the importance of selecting the most relevant metric based on dataset characteristics and performance priorities. Precision-recall curves emerge as particularly effective, especially in scenarios where correct predictions of the minority class hold greater significance. AUC-ROC or AUC-PR can be valuable when balancing true positives and false positives is crucial.

In the realm of energy consumption anomaly detection, where anomalies are rare, metrics accounting for imbalanced classes (precision, recall, F1 score, AUC-ROC, or AUC-PR) prove relevant. The continuous exploration of ensemble methods, transfer learning, and specialized architectures signals an evolving landscape, continually addressing class imbalance challenges.
Navigating the diverse perspectives presented in this review reveals the need for innovative models to robustly handle class imbalance complexities in anomaly detection for energy

consumption. This literature review lays a foundational pillar for subsequent research phases, outlining gaps, challenges, and successes that guide the development and evaluation of novel anomaly detection models tailored to the unique characteristics of imbalanced energy datasets.

4. CONCLUSION

In conclusion, the literature review provides a thorough exploration of models designed to tackle the significant challenge of class imbalance in anomaly detection for energy consumption. From classical methodologies to cutting-edge deep learning architectures, the survey illuminates the diverse array of techniques employed to enhance accuracy and reliability in anomaly detection systems for energy consumption data. The literature review emphasizes the crucial need to address class imbalance due to its pervasive impact on model performance. Through meticulous analysis, the strengths and limitations of various models are uncovered, shedding light on their adaptability to the complexities of energy consumption patterns. This research has identified several gaps in the literature landscape regarding models addressing the class imbalance challenge in energy consumption anomaly detection.

Firstly, classical machine learning models, specifically Support Vector Machines and Random Forests, still encounter challenges in effectively handling class imbalance. Future research directions should focus on improving Support Vector Machines to handle sparse anomalies scattered across the feature space, capturing nuanced patterns without solely maximizing the margin between classes. Similarly, enhancing Random Forests to detect sporadic anomalies and ensuring the collective decision of multiple trees focuses on learning the nuances of rare anomalies is critical.

Secondly, while deep learning-based models are considered state-of-the-art solutions, they also face limitations. Addressing these limitations could involve adopting a deep learning ensemble model with an effective training strategy tailored to handle class imbalance challenges in energy consumption anomaly detection.

Thirdly, although oversampling the minority class is a prevalent training method, it has limitations. Conducting an empirical study to explore various training methods and identify the most effective strategy for specific class-imbalanced datasets is crucial. These alternative training methods, successful in other domains, need empirical testing in order to evaluate their effectiveness in the context of the energy consumption anomaly detection.

Fourthly, while accuracy, AUC, F1 score, and Precision & recall are popular evaluation metrics, no single metric captures performance comprehensively. The empirical study should incorporate multiple metrics to holistically evaluate model performance.

Lastly, acknowledging differences in anomaly detection datasets, such as distribution, noise levels, and other parameters, is crucial. Conducting an empirical study on diverse datasets, considering factors like imbalance ratio, temporal aspects, feature space, types of anomalies, data volume, and labeling quality, will contribute to identifying effective training strategies and models for specific class-imbalanced datasets in energy consumption anomaly detection.

5. RECOMMENDATIONS FOR FUTURE RESEARCH

The comprehensive examination of models for addressing class imbalance challenges in energy consumption anomaly detection highlights several avenues for future research. To enhance the effectiveness of anomaly detection models, particularly in the realm of energy consumption, the following recommendations are proposed:

1. Enhancing Classical Machine Learning Models

Focus on refining Support Vector Machines to effectively handle sparse anomalies and capture nuanced patterns without solely maximizing the margin between classes. Improvement on Random Forests to detect sporadic anomalies by ensuring the collective decision of multiple trees should sufficiently focus on learning the nuances of rare anomalies.

2. Optimizing Deep Learning-Based Models

Future research should explore the adoption of state-of-the-art deep learning ensemble models with effective training strategies tailored to address class imbalance challenges in energy consumption anomaly detection.

3. Empirical Investigation on Training Methods

Future research should endeavor to conduct an empirical investigation into the performance of various state-of-the-art models, considering a diverse range of training methods. This exploration should encompass classical machine learning models, deep learning architectures, and ensemble methods. The objective would be to identify the most effective training strategy for specific classimbalanced energy consumption datasets with different distributions, levels of noise, and parameters. This exploration should include variations in imbalance ratio, temporal aspects, feature space, specific types of anomalies, data volume, and labelling quality. Understanding the impact of these diverse dataset characteristics on model performance is crucial.

4. Comprehensive Metric Evaluation

Future research should emphasize on the importance of selecting and using multiple metrics to comprehensively evaluate model performance. While accuracy, AUC, F1 score, and Precision & recall are popular, incorporating a range of metrics ensures a holistic assessment of models addressing class imbalance challenges in energy consumption anomaly detection.

5. Integration of Transfer Learning

Future research could investigate the effectiveness of transfer learning strategies in training deep learning models. This involves transferring knowledge acquired from well-balanced datasets to improve anomaly detection in imbalanced scenarios within the energy consumption domain.

6. Continued Exploration of Ensemble Methods

Future research endeavours encourage ongoing exploration of ensemble methods for deep models to mitigate overfitting challenges. The diversity provided by different models within an ensemble could help prevent overfitting and enhance overall model performance.

By delving into these research directions, future studies can contribute to the advancement of anomaly detection in energy consumption, addressing the challenges posed by class imbalance in a nuanced and effective manner.

ACKNOWLEDGEMENT

I extend my sincere gratitude to all those who contributed to the successful completion of this research endeavour. Special appreciation goes to the authors and researchers whose work laid the

foundation for the comprehensive literature review on models addressing the class imbalance challenge in energy consumption anomaly detection. Their insights provided invaluable guidance and formed the basis for the critical analysis presented in this study.

REFERENCES

[1] M. Emadaleslami and M. Haghifam, "A Machine Learning Approach to Detect Energy Fraud in Smart Distribution Network," *Int. J. Smart Electr. Eng.*, vol. 10, no. 02, Jun. 2021, doi: 10.30495/ijsee.2021.683248

[2] Z. Wu, W. Lin, and Y. Ji, "An Integrated Ensemble Learning Model for Imbalanced Fault Diagnostics and Prognostics," IEEE Access, vol. 6, pp. 8394–8402, 2018, doi: 10.1109/ACCESS.2018.2807121.

[3] R. Blagus and L. Lusa, "SMOTE for high-dimensional class-imbalanced data," *BMC Bioinformatics*, vol. 14, no. 1, p. 106, Dec. 2013, doi: 10.1186/1471-2105-14-106.

[4] K. S. Khan, R. Kunz, J. Kleijnen, and G. Antes, "Five Steps to Conducting a Systematic Review," *J. R. Soc. Med.*, vol. 96, no. 3, pp. 118–121, Mar. 2003, doi: 10.1177/014107680309600304.

[5] Z. Qu, H. Liu, Z. Wang, J. Xu, P. Zhang, and H. Zeng, "A combined genetic optimization with AdaBoost ensemble model for anomaly detection in buildings electricity consumption," Energy Build., vol. 248, p. 111193, Oct. 2021, doi: 10.1016/j.enbuild.2021.111193.

[6] N. Javaid, N. Jan, and M. U. Javed, "An adaptive synthesis to handle imbalanced big data with deep siamese network for electricity theft detection in smart grids," J. Parallel Distrib. Comput., vol. 153, pp. 44–52, Jul. 2021, doi: 10.1016/j.jpdc.2021.03.002

[7] R. ELhadad, Y.-F. Tan, and W.-N. Tan, "Anomaly Prediction in Electricity Consumption Using a Combination of Machine Learning Techniques," Int. J. Technol., vol. 13, no. 6, p. 1317, Nov. 2022, doi: 10.14716/ijtech.v13i6.5931.

[8] Rafael A., Gomez, Felipe A., and Llaugel, "Electrical Network Fraud Detection Using a CNN+LSTM Model," Int. J. Adv. Eng. Manag. Res., vol. 5, no. 8, 2023.

[9] Z. A. Khan, M. Adil, N. Javaid, M. N. Saqib, M. Shafiq, and J.-G. Choi, "Electricity Theft Detection Using Supervised Learning Techniques on Smart Meter Data," Sustainability, vol. 12, no. 19, p. 8023, Sep. 2020, doi: 10.3390/su12198023.

[10] L. J. Lepolesa, S. Achari, and L. Cheng, "Electricity Theft Detection in Smart Grids Based on Deep Neural Network," IEEE Access, vol. 10, pp. 39638–39655, 2022, doi: 10.1109/ACCESS.2022.3166146.

[11] S. K. Gunturi and D. Sarkar, "Ensemble machine learning models for the detection of energy theft," Electr. Power Syst. Res., vol. 192, p. 106904, Mar. 2021, doi: 10.1016/j.epsr.2020.106904.

[12] J. Lee, Y. C. Lee, and J. T. Kim, "Fault detection based on one-class deep learning for manufacturing applications limited to an imbalanced database," J. Manuf. Syst., vol. 57, pp. 357–366, Oct. 2020, doi: 10.1016/j.jmsy.2020.10.013.

[13] H. Qin, H. Zhou, and J. Cao, "Imbalanced learning algorithm based intelligent abnormal electricity consumption detection," Neurocomputing, vol. 402, pp. 112–123, Aug. 2020, doi: 10.1016/j.neucom.2020.03.085.

[14] L. Zhang, X. Chen, C. Zhang, L. Zhang, and B. Zou, "Intelligent Anomaly Detection Method of Gateway Electrical Energy Metering Devices using Deep Learning," Int. J. Adv. Comput. Sci. Appl., vol. 14, no. 7, 2023, doi: 10.14569/IJACSA.2023.0140793.

[15] A. Castellani, S. Schmitt, and S. Squartini, "Real-World Anomaly Detection by Using Digital Twin Systems and Weakly Supervised Learning," IEEE Trans. Ind. Inform., vol. 17, no. 7, pp. 4733–4742, Jul. 2021, doi: 10.1109/TII.2020.3019788.

[16] D. Cao, D. Liu, X. Ren, and N. Ma, "Self-Adaption AAE-GAN for Aluminum Electrolytic Cell Anomaly Detection," IEEE Access, vol. 9, pp. 100991–101002, 2021, doi: 10.1109/ACCESS.2021.3097116

AUTHORS

David Kaimenyi Marangu is a Senior Technologist in the Department of Information Technology, School of Computing and Information Technology, Murang'a University of Technology, Kenya. He earned his Bachelor of Education (Science) degree from Kabarak University, MSc in Information Technology from Moi University, and MSc in Data Analytics from KCA University. He is currently pursuing his PhD in Information Technology at Murang'a University of Technology. His research interests are in Artificial intelligence and data analytics,

Dr Stephen Thiiru Njenga is a lecturer in the Department of Computer Science, School of Computing and Information Technology, Murang'a University of Technology, Kenya. He has more than fifteen years of teaching at the University level. His areas of research include Machine Learning, Intelligent Agents and Distributed Ledger Technology.

Dr. Rachael Njeri Ndung'u is a Lecturer at the Department of Information Technology, School of Computing and Information Technology, Murang'a University of Technology, Kenya. She holds a Ph.D. in Information Technology. Her research interests include Artificial Intelligence, Machine Learning, Data Analytics, and Blockchain.