

# Explaining deep learning-based anomaly detection in energy consumption data by focusing on contextually relevant data

Mohammad Noorchenarboo, Katarina Grolinger<sup>1</sup>,\*

Department of Electrical and Computer Engineering, Western University, 1151 Richmond St, London, N6A 3K7, Ontario, Canada

## ARTICLE INFO

### Keywords:

Explainable Artificial Intelligence (XAI)  
SHAP  
Anomaly detection  
Energy consumption  
Deep learning  
Baseline  
Background dataset  
Model explainability

## ABSTRACT

Detecting anomalies in energy consumption data is crucial for identifying energy waste, equipment malfunction, and overall, for ensuring efficient energy management. Machine learning, and specifically deep learning approaches, have been greatly successful in anomaly detection; however, they are black-box approaches that do not provide transparency or explanations. SHAP and its variants have been proposed to explain these models, but they suffer from high computational complexity (SHAP) or instability and inconsistency (e.g., Kernel SHAP). To address these challenges, this paper proposes an explainability approach for anomalies in energy consumption data that focuses on context-relevant information. The proposed approach leverages existing explainability techniques, focusing on SHAP variants, together with global feature importance and weighted cosine similarity to select background dataset based on the context of each anomaly point. By focusing on the context and most relevant features, this approach mitigates the instability of explainability algorithms. Experimental results across 10 different machine learning models, five datasets, and five XAI techniques, demonstrate that our method reduces the variability of explanations providing consistent explanations. Statistical analyses confirm the robustness of our approach, showing an average reduction in variability of approximately 38% across multiple datasets.

## 1. Introduction

The rising global demand for energy, particularly electricity, is causing significant environmental impacts, including increased greenhouse gas emissions and the depletion of vital resources [1]. Global electricity demand is anticipated to rise by nearly 80% by 2040 [2]. Residential and commercial buildings, accounting for one-third of global energy consumption, are major contributors to these environmental impacts [1]. Improving building energy efficiency, especially in terms of electricity usage, is essential to mitigate the adverse effects of growing energy consumption [3]. Achieving this goal relies heavily on detecting and correcting anomalies in electricity consumption. These anomalies, defined as irregularities or deviations from normal energy behavior, include unusual consumption patterns caused by faulty device operations, user negligence (e.g., leaving windows or refrigerator doors open), theft, or non-technical losses. If not addressed promptly, these issues can result in energy waste, increased power consumption, and devices running longer than necessary due to inefficiencies or malfunctions, leading to additional energy waste and potential equipment damage [4–6].

Data-driven approaches have proven effective in identifying anomalies, offering reliable alerts to analysts and energy managers [7,8]. Anomaly detection has evolved from traditional statistical methods, which struggle with complex structures and large datasets, to Deep Learning (DL) methods that automatically learn from time series data [9–11]. DL algorithms excel in identifying abnormal electricity consumption patterns due to their ability to model complex non-linear relationships and leverage multi-layered architectures for hierarchical feature extraction [12,13]. Despite their superior accuracy, DL models face challenges in transparency and explainability, crucial factors for building trust and ensuring successful real-world deployment [14]. Providing clear, instance-specific explanations for anomalies are needed to enhance expert trust, support informed decisions, and facilitate the adoption of complex DL models in the energy sector [15].

Explainable Artificial Intelligence (XAI) aims to enhance transparency and provide explanations, allowing users to understand and trust them [16]. While XAI methods have been primarily used in computer vision to provide visual interpretability and explain decisions in tasks such as object recognition, they are equally important for time se-

\* Corresponding author.

E-mail address: [kgroling@uwo.ca](mailto:kgroling@uwo.ca) (K. Grolinger).

<https://doi.org/10.1016/j.enbuild.2024.115177>

Received 21 October 2024; Received in revised form 5 December 2024; Accepted 9 December 2024

List of Acronyms and Symbols

Acronym/Symbol	Definition
1D-CNN	One-dimensional Convolutional Neural Network
AI	Artificial Intelligence
BaseValue	Average prediction made using the background dataset
BGRU	Bidirectional Gated Recurrent Unit
BLSTM	Bidirectional Long Short-Term Memory
CNN	Convolutional Neural Network
DCNN	Dilated Convolutional Neural Network
DL	Deep Learning
DNN	Deep Neural Network
$e$	Prediction error
$ F $	Total number of features in the model
$GFI_i$	Global Feature Importance score for feature $i$
$GFI'_i$	Exponentially transformed Global Feature Importance score for feature $i$
GRU	Gated Recurrent Unit
$h$	Forecasting horizon
$I$	Number of time steps for inputs sequence
$IQR$	Interquartile range, calculated as $Q_3 - Q_1$
$k(F, S)$	Kernel weights ensuring fair evaluation of feature contributions
LIME	Local Interpretable Model-Agnostic Explanations
LSTM	Long Short-Term Memory
$O_i$	Output sequence of target values in horizon $i$
$Q_1$	First quartile (25th percentile) of the prediction errors
$Q_3$	Third quartile (75th percentile) of the prediction errors
RNN	Recurrent Neural Network
$S(x_c, x_a)$	Weighted cosine similarity score between $x_c$ and $x_a$
SD	Standard Deviation
$ S $	Number of features in subset $S$
SHAP	Shapley Additive Explanations
TCN	Temporal Convolutional Network
TFT	Temporal Fusion Transformer
TPE	Tree-structured Parzen Estimator
TST	Time Series Transformer
$v(S)$	Model value function evaluated for subset $S$
WLS	Weighted Least Squares
XAI	Explainable Artificial Intelligence
$X$	Original feature values
$X'$	Scaled feature values
$X_{f,t}$	Value of feature $f$ at time $t$
$X_{\max}$	Maximum value of the feature
$X_{\min}$	Minimum value of the feature
$x_{ai}$	Value of feature $i$ for the anomalous sample
$x_{ci}$	Value of feature $i$ for a sample from the training dataset
$Y_{t+h}$	Target energy consumption values for time steps $t+h$
$\binom{ F }{ S }$	Binomial coefficient representing the number of ways to choose $ S $ features from $ F $
$\phi_0$	Base value representing the average prediction over the entire background dataset
$\phi_i$	SHAP value representing feature $i$ 's contribution to the prediction

ries data to understand decision-making processes. As sensors become more affordable and ubiquitous, generating vast amounts of electricity consumption time series data, analysis of these data can automate tasks such as energy usage monitoring to enhance maintenance and reduce inefficiencies. For DL-based time series models, such as Long Short-Term Memory (LSTM) networks or Transformers, data are typically transformed into time-series segments using the sliding window technique, necessitating the adaptation of current explainability approaches to provide meaningful explanations for electricity consumption analysis based on this segmented data structure. Converting time points into features for local feature importance and visualizing relevance, similar to saliency masks in images, will provide experts with insights into the decision-making process and facilitate energy improvement tasks [17].

While XAI approaches are generally categorized as either model-specific, tailored to particular model architectures, or model-agnostic, applicable to any model by focusing on input-output relationships, this study focuses on model-agnostic methods due to their versatility and broad applicability across diverse machine learning models. A prominent model-agnostic approach is SHapley Additive exPlanations (SHAP), which is widely recognized for its ability to provide consistent and interpretable feature importance scores using Shapley values from game theory [18]. SHAP has gained considerable attention due to its solid theoretical foundation and its effectiveness in delivering reliable expla-

nations across various domains, making it a widely accepted method for local explanations [19].

Despite the great success of SHAP, its application to large datasets is often hindered by its computational complexity. Calculating exact SHAP values is not only time-consuming but can also be impractical for many widely used models. To address these challenges, approximation methods such as Kernel SHAP have been proposed. Kernel SHAP estimates SHAP values by solving a weighted linear regression on a sample of perturbed instances, providing a more computationally efficient means of obtaining SHAP value estimates with fewer evaluations of the original model [20].

However, while Kernel SHAP offers improvements in efficiency, it introduces a new challenge: instability in explanations [21]. Kernel SHAP, as well as its variants Partitioning, and Sampling, depends on the choice of a background dataset as the baseline, which serves as a reference point for SHAP explanations and can strongly impact feature attribution values. Different executions of Kernel SHAP with the same inputs can produce varying explanations, leading to inconsistencies that can undermine the reliability of the results and decrease user trust [21,22], particularly in anomaly detection for energy consumption data, where stable and consistent explanations are essential for effective decision-making and energy management.

To address these challenges, this paper proposes an approach for explaining anomaly detection models in energy consumption data that

mitigates instabilities and enhances explanation reliability, particularly in SHAP-based techniques. The consistency of explanations is improved by selecting data samples relevant to the anomaly under observation when determining the baseline for SHAP value calculations. Moreover, our approach prioritizes the features most relevant to anomalies, providing consistent and accurate insights while minimizing computational burden by selecting only a few background data samples. Evaluation results across five datasets and five XAI techniques demonstrate improvements in explanation stability and reliability, with reduced variability in feature importance.

The rest of this paper is structured as follows: Section 2 reviews related work, Section 3 provides the background, Section 4 details the methodology, Section 5 presents the results and discusses the finding, and Section 6 concludes the study.

## 2. Related work

This section reviews XAI techniques in energy systems and time series applications, as well as in anomaly and fault detection.

### 2.1. XAI in energy systems and time series applications

Several studies have leveraged XAI to provide insights into workings of energy load forecasting models and other building energy management systems. Moon et al. [23] proposed Explainable Electrical Load Forecasting (XELF) methodology for educational buildings, emphasizing the importance of understanding and interpreting the factors that influence electrical load predictions. By incorporating external factors such as weather and internal building data, they trained various tree-based models and utilized SHAP to provide interpretable explanations for the energy predictions made by these models. Similarly, Chung and Liu [24] analyzed input variables for deep learning models predicting building energy loads, comparing the XAI techniques Local Interpretable Model-Agnostic Explanations (LIME) and SHAP, and found that SHAP outperformed LIME by maintaining prediction accuracy with fewer input variables. Joshi et al. [25] presented a data-driven approach for benchmarking energy usage in Singapore, employing ensemble tree models and XAI techniques such as SHAP for a detailed analysis of the impact of building attributes on energy consumption.

The integration of XAI in power systems and renewable energy has also been explored. Zhang et al. [26] utilized SHAP to explain deep reinforcement learning models for power system emergency control, generating SHAP values to quantify the impact of each system variable and clarify how different factors influenced emergency control decisions. Tan et al. [27] proposed an explainable Bayesian neural network for probabilistic transient stability analysis in power systems, using the Gradient SHAP algorithm for explanations. Their approach provided insights at both global and local levels, with global explanations offering a comprehensive understanding of factors influencing overall model behavior and local explanations detailing individual predictions.

Leuthe et al. [28] explored XAI in building energy consumption forecasting and compared transparent and black-box models. They considered linear regression, decision tree, and QLattice as transparent prediction models and applied four XAI methods - partial dependency plots, Accumulated Local Effects (ALE), LIME, and SHAP - to an artificial neural network using a real-world dataset of residential buildings. Their findings indicate that appropriate XAI methods can significantly improve decision-makers' satisfaction and trust in machine learning models for energy forecasting. Mueller et al. [29] examined the use of XAI to explain characteristics of vehicle power consumption, referring to the energy consumption within a vehicle's low-voltage electrical system. The study applied methods such as ALE and Permutation Feature Importance (PFI) for global insights, and LIME and SHAP for local analysis.

XAI techniques have also been investigated in time series analysis beyond the energy domain. Rožanec et al. [30] examined the ex-

plainability of global time series forecasting models. They used the M4 competition and Kaggle Wikipedia Web Traffic datasets, integrating anomaly detection with XAI techniques to identify and explain poor forecasts. Anomaly detection flagged deviations, while XAI methods, such as LIME, computed feature attributions and generated counterfactual examples to elucidate the reasons behind these deviations. In addition, Labaien Soto et al. [31] proposed a model-agnostic approach that uses autoencoders to generate real-time counterfactual explanations. This approach used one-dimensional Convolutional Neural Networks (1D-CNN) and Recurrent Neural Networks (RNNs) to analyze time series vibration data. The approach explains anomalies by making small input modifications to measure how far they deviate from normal behavior, providing insights into the factors contributing to the anomaly.

Freeborough and van Zyl [32] explored XAI methods for financial time series forecasting, focusing on ablation, integrated gradients, added noise, and permutation techniques to assess feature importance and enhance model explainability. These methods were applied with several ML models including standard RNN, LSTM, and Gated Recurrent Unit (GRU). Schlegel and Keim [33] focused on enhancing the explainability of deep learning models for time series data using perturbation analysis and evaluated state-of-the-art XAI techniques, including gradient-based methods (Saliency, Integrated Gradients) and SHAP-based methods (DeepLiftShap, Kernel SHAP), on three time-series classification datasets: FordA, FordB, and ElectricDevices. The study analyzed attribution distributions, skewness, and both Euclidean and cosine distances between original and perturbed instances. It found that SHAP and its derivatives generally produce effective attributions, while noting that the efficacy of XAI techniques can vary significantly depending on the chosen perturbation strategy.

Despite significant advancements in using XAI techniques for energy load forecasting and time series analysis, current methods do not adequately address the reliability of explanation results. Given that SHAP methods are grounded in tabular data, they needed to be adapted to accommodate time series data. SHAP struggles with computation complexity while SHAP variants result in unstable explanations. To address this, we propose a technique leveraging SHAP approaches, but address instability through targeted selection of the background dataset based on context-relevant information.

### 2.2. XAI in anomaly/fault detection

XAI has been instrumental for enhancing the explainability of anomaly detection and fault detection models. Roshan and Zafar [34] explored the use of SHAP for feature selection in an unsupervised anomaly detection setting. Their approach leveraged SHAP to improve the performance of autoencoders by identifying key features responsible for anomalies, and retraining the model using only benign data. While the SHAP-based feature selection showed improved results over other methods, the paper faced limitations in computational cost due to Kernel SHAP's complexity and potential sampling bias from using a subset of the CICIDS2017 dataset, which could affect generalizability. On the other hand, Antwarg et al. [15] focused on applying the SHAP framework, traditionally used in supervised learning, to explain anomalies detected by unsupervised autoencoders. Their approach emphasizes understanding the relationships between features with high reconstruction errors and those most critical to anomaly detection. Kernel SHAP is utilized to calculate feature importance, offering detailed insights into why certain anomalies occur, by identifying both contributing and offsetting features. However, challenges such as selecting the appropriate background dataset and further validation across various autoencoder architectures remain areas for future research. While Roshan and Zafar's work primarily aimed at enhancing model prediction accuracy through feature selection, Antwarg et al. focused more on interpreting feature contributions and relationships in the context of anomaly explanations.

Kim et al. [35] proposed an explainable anomaly detection framework for maritime main engine sensor data, combining SHAP with hier-

archical clustering to interpret and group common anomaly patterns. By transforming SHAP values based on their distributions, Kim et al. were able to identify and isolate key sensor variables contributing to anomalies, allowing for more precise segmentation and analysis of the detected anomalies. This method provides insights into the causes of anomalies by visualizing and grouping similar patterns, offering an improved understanding of the conditions leading to engine failures. A two-layer patient monitoring system employing Kernel SHAP for anomaly detection and explanation in healthcare data was presented by Abououf et al. [36]. The most influential features contributing to anomalies detected by an autoencoder-based model were identified utilizing SHAP.

Kernel SHAP was also used for fault detection in industrial applications. Asutkar and Tallur [37] proposed an explainable unsupervised learning framework using autoencoders for fault detection and Kernel SHAP for explanations. Their framework accurately identified machine faults under varying operating conditions, with Kernel SHAP highlighting the most prominent features contributing to fault detection. The approach was validated using multiple datasets, demonstrating its effectiveness and scalability in real-world industrial applications. For process monitoring, Choi and Lee [38] proposed an explainable fault diagnosis model that combines Stacked Autoencoders (SAE) with Kernel SHAP for feature importance and model behavior explanation. Kernel SHAP provided a clear explanation of which features contributed most significantly to the model's predictions.

The Table 1 summarizes key studies that have applied XAI techniques, especially SHAP variants, to energy consumption and related fields. These works highlight the benefits of using XAI methods to improve model transparency, interpretability, and decision-maker trust. For example, SHAP methods have been effectively used to analyze energy load forecasting, identify influential features, and provide both global and local explanations for model predictions. By incorporating XAI techniques, these studies have enhanced the understanding of how external factors such as weather, building attributes, and sensor data influence predictions, enabling more informed decision-making in energy systems.

Although XAI methods such as SHAP have been successfully applied to anomaly detection and fault diagnosis, particularly with autoencoders, a gap remains in ensuring the reliability and stability of explanations. Current approaches suffer from inconsistencies, as noted in multiple studies [34,15], due to the random sampling of training data for the background dataset. For instance, Chung and Liu [24] noted variability in SHAP explanations due to random background data selection and input dependencies, while Antwarg et al. [15] and Chen et al. [22] highlighted challenges in selecting appropriate background datasets, which significantly influences explanation stability. Similarly, Schlegel and Keim [33] reported inconsistencies tied to perturbation strategies, often exacerbated by arbitrary baseline datasets, and Roshan and Zafar [34] documented issues with sampling bias impacting explanation robustness. We address these issues by proposing a robust approach that provides stable explanations focusing on highly relevant data for the anomaly under consideration. By aligning baselines with the contextual data of anomalies, our method provides consistent and reliable insights, representing a significant improvement over previous practices and enhancing the practical utility of explainability in anomaly detection.

### 3. Background

This section first introduces core concepts in explaining black box models. Next, classical SHAP and Kernel SHAP are introduced, as our approach leverages these techniques.

#### 3.1. Explaining black-box models

Understanding and interpreting the decisions made by Artificial Intelligence (AI) models, especially those that function as black-box systems, is crucial for ensuring their legitimacy and reliability in sensitive

applications [39]. Modern AI models such as Deep Neural Networks (DNNs) are complex systems with many parameters, making them difficult to interpret [40]. To address this challenge XAI has emerged, offering two key types of explainability: local and global. Local explainability focuses on explaining the decision for a specific instance, offering detailed insights into how a particular prediction was made. In contrast, global explainability provides an overarching view of the model's decision-making process, giving a broader understanding of the AI system's behavior across its entire input space [41,42]. Both local and global explainability play important roles in making AI systems more transparent, accountable, and understandable, thereby enhancing trust and facilitating decision-making [42].

In this research, we focus on local explainability because it allows for a deeper understanding of individual anomalies detected within the data, assisting in determining the root cause of each anomaly. Both LIME and SHAP provide local explanations. LIME explains one prediction at a time by constructing a simple linear model around the data point, using random perturbation to create simulated data. However, this can lead to instability of explanations. Conversely, SHAP provides individual explanations by assigning feature importance based on Shapley values, offering a more stable approach [43].

#### 3.2. Classic SHAP and Kernel SHAP

Shapley value estimation, based on cooperative game theory, calculates the contribution of each feature by comparing the model's predictions with and without that feature across all possible feature combinations [18]. However, this is computationally expensive, especially with large feature sets. To address this, sampling methods approximate Shapley values without requiring retraining for every combination of features [44]. Despite these approximations, the process remains resource-intensive for large datasets, leading to the development of more efficient methods such as Kernel SHAP.

Building on the classic Shapley value estimation, Kernel SHAP provides a more practical approach to approximate Shapley values by solving a Weighted Least Squares (WLS) problem. This weighted approach assigns different importance levels to each possible subset of features. Kernel SHAP utilizes sampling techniques to approximate Shapley values while reducing computational burden, thus making it feasible for models with many features while maintaining their theoretical integrity [18,21]. This balance of practicality and rigor makes Kernel SHAP well suited for energy applications; nevertheless, it still suffers from instability. The Kernel SHAP estimates the contribution of each feature through the following WLS problem:

$$\min_{\phi_0, \dots, \phi_F} \sum_{S \subseteq F} k(F, S) \left( v(S) - \left( \phi_0 + \sum_{j \in S} \phi_j \right) \right)^2 \quad (1)$$

Here  $F$  represents the total number of features in the model, and  $S$  denotes a subset of features. The term  $v(S)$  refers to the value of the predictive model when only the features in subset  $S$  are considered. The term  $\phi_0$  is the base value representing the average prediction over the entire background dataset. The contribution of feature  $j$  indicated by  $\phi_j$  is calculated as part of the WLS optimization process by balancing the error terms across all feature subsets. The weights  $k(F, S)$  in Equation 1, are given by:

$$k(F, S) = \frac{(|F| - 1)}{\binom{|F|}{|S|} \cdot |S| \cdot (|F| - |S|)} \quad (2)$$

Here  $\binom{|F|}{|S|}$  is the binomial coefficient representing the number of ways to choose  $|S|$  features from the total  $|F|$  features. A subset  $S$  refers to any possible combination of features from the full set  $F$ . For example, given three features  $\{A, B, C\}$ , possible subsets include  $\{A\}$ ,  $\{B, C\}$ ,  $\{A, B, C\}$ , and the empty set  $\{\}$ . The kernel weights  $k(F, S)$ , defined in Equation 2, are derived from cooperative game theory, where the contribution of each feature is evaluated by considering

**Table 1**  
Summary of related works.

Reference	Main Points	Pros	Cons
Moon et al. (2022)	Proposed XELF methodology for educational buildings using SHAP to interpret energy load predictions.	Improved transparency and accuracy.	Focused on educational buildings only, limited datasets.
Chung and Liu (2022)	Compared LIME and SHAP for building load prediction: SHAP reduced variables while maintaining accuracy.	Reduced input features, maintained accuracy, improved interpretability.	Limited to LIME and SHAP comparison, excluding other methods.
Joshi et al. (2022)	Benchmarked energy in Singapore buildings using ensemble trees and SHAP to analyze attributes.	Real data, enhanced interpretability, and benchmarking insights.	Limited to Singapore hotel and retail buildings.
Zhang et al. (2020)	Used SHAP to explain deep reinforcement learning models for power system emergency control.	Quantified variable impact, improved decision transparency.	Focused on power systems, high complexity of variable interpretation.
Tan et al. (2022)	An explainable Bayesian NN for probabilistic transient stability analysis in power systems using Gradient SHAP.	Global and local insights.	Lacks comparison with alternative methods and validation of result reliability.
Leuthe et al. (2024)	Explored XAI in building energy forecasting, applying XAI methods to ANN models with residential datasets.	Enhanced trust with transparent models, compared XAI techniques.	Limited to a single dataset, lacks diversity in building types.
Mueller et al. (2023)	Applied XAI methods (ALE, PFI, LIME, SHAP) to vehicle power consumption in low-voltage systems.	Combined global/local insights, enhanced transparency for complex systems.	Lacks comparison of reliability between global and local insights.
Rožanec et al. (2021)	Combined anomaly detection and XAI for time series forecasting, provided counterfactual explanations for poor forecasts.	Identified deviations, clarified anomalies with counterfactuals.	Lacks evaluation of broader explainability approaches.
Labaien Soto et al. (2023)	A model-agnostic approach with autoencoders for counterfactual explanations in anomaly detection.	Real-time explanations, works across model types.	Lacks comparison to diverse deep learning architectures.
Freeborough and van Zyl (2022)	XAI methods (e.g., ablation, integrated gradients, permutation) for financial time series with RNNs.	Effective feature attribution for RNNs, LSTMs, and GRUs.	Lacks assessment of XAI reliability and comparison across scenarios.
Schlegel and Keim (2023)	Deep learning explainability for time series using gradient- and SHAP-based methods across multiple datasets.	Provided detailed attribution insights, effective perturbation analysis.	Limited robustness evaluation and consistency analysis of XAI methods.
Roshan and Zafar (2021)	SHAP for feature selection in unsupervised anomaly detection with autoencoders.	Improved anomaly detection accuracy, highlighted key features.	High computational, sampling bias in datasets.
Antwarg et al. (2021)	SHAP to explain anomalies detected by unsupervised autoencoders, analyzing relationships between features.	Detailed insights into feature contributions and relationships.	Explanation stability and generalization across scenarios are unexamined.
Kim et al. (2021)	Developed an explainable anomaly detection framework for maritime engine data combining SHAP and clustering.	Identified key variables, segmented anomaly patterns.	Does not evaluate the robustness of explanations across anomaly patterns.
Abououf et al. (2023)	Proposed a two-layer patient monitoring system using Kernel SHAP for anomaly detection in healthcare data.	Improved anomaly interpretability for healthcare monitoring.	High computational cost, limited to specific healthcare applications.
Asutkar and Tallur (2023)	Used Kernel SHAP in an unsupervised framework for fault detection in industrial machines.	Effective under varying conditions, scalable for Industry 4.0.	Limited exploration of explainability in multi-fault scenarios.
Choi and Lee (2022)	Combined SAE and Kernel SHAP for explainable fault diagnosis in process monitoring.	Clear feature importance, high classification accuracy.	Limited to specific industrial process data.

all possible subsets  $S$  of features. These weights  $k(F, S)$  ensure that each feature's contribution is assessed in a balanced manner by giving different levels of importance to subsets of different sizes. More specifically, the weighting ensures that features are not biased by their position in a subset and that smaller subsets are given fair consideration [21]. This weighting scheme balances the evaluation, ensuring that each feature's contribution is assessed within a meaningful and fair context.

Both Classic SHAP and Kernel SHAP determine feature importance by excluding features from the model to observe their impact on predictions. Since the model is already trained and features cannot be physically removed, features are substituted with alternative values to reduce their influence. This substitution typically involves using values

from a designated background dataset. When calculating the model's prediction without a specific feature, the real value of that feature is replaced with a value from the background dataset. The background dataset can consist of either the entire training data or, in the case of large datasets, a representative subset. This allows the model to simulate the absence of the feature and assess its contribution by comparing changes in predictions when the feature is excluded.

#### 4. Methodology

This section presents the proposed approach for explaining the anomaly detection model for energy consumption data by leveraging

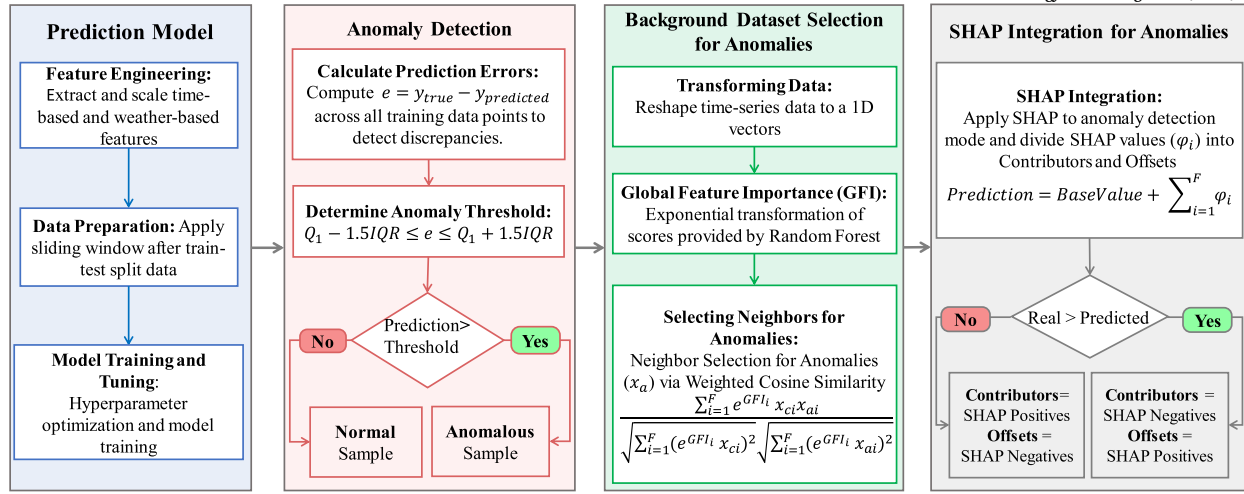


Fig. 1. Methodology overview.

variants of SHAP, such as Kernel SHAP, but improving stability and consistency through the targeted selection of the background dataset (baseline) using a weighted cosine similarity technique. The approach is specifically designed for prediction-based anomaly detection techniques where a black-box model is used to generate energy predictions, which are in turn compared to actual energy values. If the difference exceeds the threshold, the sample is deemed anomalous. After detecting anomalies, we proceed to the explanation phase, which involves multiple steps to provide clear feature contributions for each anomaly. The overview of the complete process of detecting anomalies and explaining underlying features is presented in Fig. 1, while the details of each component are presented in the following subsections.

#### 4.1. Prediction model

This first step involves getting the prediction-based model ready for the anomaly detection. As seen from Fig. 1, it consists of feature engineering, data preparation, and model training and tuning.

##### 4.1.1. Feature engineering

Energy consumption data obtained from smart meters or other sensors typically consists of energy consumption and the reading date-time recorded in hourly or similar intervals. For better anomaly detection, we extract the following features from the reading date-time: the hour of the day, the day of the week, the day of the month, the day of the year, the month, and an indicator for weekends. Weather-related features such as temperature, humidity, and wind speed are also incorporated, along with previous energy consumption readings as input features, while the energy consumption remains the target variable. Data are scaled using Min-Max scaling to bring all features to a similar range and avoid dominance of large numbers:

$$X' = \frac{X - X_{\min}}{X_{\max} - X_{\min}} \quad (3)$$

where  $X$  represents the original feature values,  $X_{\min}$  and  $X_{\max}$  are the minimum and maximum values of the feature, and  $X'$  are the scaled values.

##### 4.1.2. Data preparation

Next, the dataset is chronologically split, with the first 80% used for training, the next 10% for validation, and the final 10% for testing. To capture temporal dependencies, the sliding window technique is applied, moving a fixed-length window along the time series, advancing one record at a time. As shown in Fig. 2, this approach creates sequences of inputs and corresponding outputs for the model to learn temporal re-

lationships. For each window  $i$ , the input sequence is represented as a matrix of time steps and features:

$$\begin{bmatrix} X_{1,t_1} & X_{2,t_1} & \dots & X_{F,t_1} \\ X_{1,t_2} & X_{2,t_2} & \dots & X_{F,t_2} \\ \vdots & \vdots & \ddots & \vdots \\ X_{1,t_I} & X_{2,t_I} & \dots & X_{F,t_I} \end{bmatrix} \quad (4)$$

where  $I$  represents the number of time steps in the sliding window,  $F$  is the number of features, and  $X_{f,t}$  is the value of feature  $f$  at time  $t$ .

The output sequence for window  $i$  is represented as:

$$O_i = [Y_{t+1}, Y_{t+2}, \dots, Y_{t+h}] \quad (5)$$

where  $h$  is the forecasting horizon and  $[Y_{t+1}, Y_{t+2}, \dots, Y_{t+h}]$  are the target energy consumption values for time steps  $t+1$  to  $t+h$ . In this paper, we used four time steps ( $I = 48$ ) for inputs and one step ahead forecasting horizon ( $h = 24$ ) as shown in Fig. 2, but explanation only provided for first horizon ( $h = 1$ ). The prediction model supports anomaly detection through comparison of predicted and actual values.

##### 4.1.3. Model training and tuning

Here, the prediction model, regardless of which type of architecture is employed, is trained and tuned. The selection of hyperparameters that need to be tuned depends on the selected model, but the tuning process remains the same. Bayesian optimization, specifically using the Tree-structured Parzen Estimator (TPE) [45], is employed for tuning due to its resource efficiency, although other techniques could be used as well. Tuning is carried out with the validation set.

The optimization process starts with defining the hyperparameter search space, outlining the ranges and types of hyperparameters to be optimized. The TPE algorithm is then initialized, and the model performance is assessed based on the initial hyperparameters. After each evaluation, the model updates the hyperparameter probabilities, and new, more promising hyperparameters are selected. This cycle of evaluation and refinement continues until the optimal set of hyperparameters is found. Once hyperparameters are selected through TPE, the model is trained using these hyperparameters on the training data. Finally, the trained prediction model is ready for anomaly detection.

#### 4.2. Anomaly detection

As already mentioned, prediction-based anomaly detection methods, including those employed in this study, identify the anomalies by comparing the predicted values with the actual values. Normal data is expected to have small deviations, while anomalous samples are expected to result in larger discrepancies. As seen from Fig. 1, anomaly




Fig. 2. Illustration of the sliding window approach for sequence-to-sequence model training with 48 Input sequence and 24 output sequence.




Fig. 3. Illustration of the anomaly detection process using IQR.

detection involves calculating prediction error, determining an anomaly threshold, and classifying samples as anomalous or non-anomalous.

Next, we calculate the *prediction error* for the entire training dataset. The prediction error is defined as the difference between the actual and predicted values for each data point within a single prediction window. By computing the prediction errors across all data points in the training dataset, we obtain a comprehensive distribution of errors.

Next, the anomaly threshold is determined using the Interquartile Range (IQR) method calculated based on the prediction errors from all training data, as it is well-suited for identifying outliers in skewed data distributions [46]. The IQR method sets a range that defines normal data behavior, as shown in Equation 6:

$$Q_1 - 1.5 \times IQR \leq e \leq Q_3 + 1.5 \times IQR \quad (6)$$

where  $Q_1$  represents the first quartile (25th percentile) of the prediction errors,  $Q_3$  is the third quartile (75th percentile), and  $IQR$  is the interquartile range, calculated as  $Q_3 - Q_1$ . Fig. 3 illustrates this anomaly detection process using the IQR method.

Each prediction error is compared against the established threshold. If the prediction error exceeds this threshold, the data point is classified as an *Anomalous Sample*, suggesting that it deviates significantly from expected behavior. Conversely, if the prediction error falls within the threshold, the data point is classified as a *Probably Normal Sample*. This classification enables the identification of potential issues in energy consumption.

### 4.3. Background dataset selection for anomalies

In the context of SHAP value explanations, the background dataset serves as the baseline, representing the expected or average behavior against which individual predictions are compared. According to Chen et al. [22], selecting an appropriate baseline is crucial as it defines how absent feature values are handled during Shapley value calculations. Different baseline strategies, such as fixed baselines or distributional baselines (marginal and conditional), influence the resulting SHAP val-

ues in distinct ways. Our approach of selecting a similar background dataset using weighted cosine similarity aligns with the concept of distributional baselines, where the background samples are chosen to reflect the statistical context of the anomalies being explained. By carefully selecting neighbors that closely resemble the anomaly points, we ensure that the baseline accurately captures the relevant feature distributions and dependencies, thereby enhancing the interpretability and reliability of the SHAP value explanations. This method mitigates the arbitrariness associated with fixed baselines and preserves the integrity of feature relationships, particularly in complex, high-dimensional datasets. The process is presented in Algorithm 1 and detailed in the following subsections.

#### 4.3.1. Transforming data

As our approach leverages the SHAP algorithm, it is necessary to first transform the windowed time series data into a format compatible with this method, which operates on tabular datasets with row vectors as samples. This transformation process is the first step of Algorithm 1. For anomaly detection, the energy data is represented in a windowed format as shown in Equation 4. These data are transformed into a tabular format  $(1 \times I \times F)$ , representing the multi-dimensional input into a one-dimensional vector. The transformed data are expressed as follows:

$$[X_{11}, X_{12}, \dots, X_{1F}, X_{21}, X_{22}, \dots, X_{2F}, \dots, X_{IF}]_{(1 \times I \times F)} \quad (7)$$

For example, when the model uses inputs with a window length of 48 and 10 features, this transformation results in a 480-dimensional vector representing the input, with explanations provided for each feature.

This paper focuses on explaining one-step-ahead predictions, corresponding to the first horizon of the output sequence  $O_i = [Y_{i+1}]$ . By setting the output window  $h$  to 1, this approach facilitates early anomaly detection. However, the same approach can be extended to explain predictions for multiple horizons. To do this, one only needs to substitute the first horizon with the desired horizon (e.g., the second horizon) and repeat the same transformation and evaluation steps to provide explanations for subsequent horizons.

#### 4.3.2. Global feature importance

After data are transformed into one-dimensional vectors, a surrogate model is used to calculate global feature importance (Step 2 of Algorithm 1). While various models can serve this purpose, we selected Random Forest due to its speed and computational effectiveness. Instead of using approaches that focus primarily on local explanations, such as SHAP, we employed Random Forest to provide faster global feature importance estimates. This is because it uses Gini importance,

which does not require calculating local feature importance for each individual training data point. This advantage makes it well-suited for large datasets while maintaining high accuracy [47]. Moreover, Random Forest's ability to handle high-dimensional datasets and robust performance across various datasets made it an excellent choice for this anomaly explanation step [48].

To further strengthen the feature importance derived from the Random Forest model, the exponential transformation is applied to the importance scores, defined as  $GFI'_i = \exp(GFI_i)$ . This transformation increases the differences in importance among features, giving more important features greater impact in subsequent analyses. As a result, the exponential scaling makes similarity measures, specifically cosine similarity, pay more attention to important features, improving the selection of relevant neighbors for SHAP explanations.

#### 4.3.3. Selecting neighbors for anomalies

In step 3 of Algorithm 1, weighted K-Nearest Neighbor (KNN) is employed to select neighbors for each detected anomaly point under analysis. The similarity score between an anomaly point  $x_a$  and each point  $x_c$  in the training set is computed using the weighted cosine similarity metric:

$$S(x_c, x_a) = \frac{\sum_{i=1}^F GFI'_i \cdot x_{ci} \cdot x_{ai}}{\sqrt{\sum_{i=1}^F (GFI'_i \cdot x_{ci})^2} \cdot \sqrt{\sum_{i=1}^F (GFI'_i \cdot x_{ai})^2}} \quad (8)$$

Here,  $x_{ai}$  and  $x_{ci}$  represent the values of feature  $i$  for the anomalous sample and another sample from the training dataset, respectively, and  $F$  denotes the total number of features. The weight of feature  $i$ ,  $GFI'_i$ , is the global feature importance calculated determined as described in Subsection 4.3.2. For each anomaly point, the top 100 most similar samples are initially selected as the background dataset for calculating local SHAP values. Future work may explore the impact of varying the number of selected samples on both computational efficiency and the quality of SHAP explanations to further optimize the process.

---

#### Algorithm 1 Background dataset selection for anomaly explanation.

---

- 1: **Input:** Training data  $X_{\text{train}}$ , Test data  $X_{\text{test}}$ , Trained model  $M$
  - 2: **Output:** Background dataset  $B$  for anomaly explanation
  - 3: **Step 1: Transform Windowed Data into 1D Format**
  - 4: Reshape windowed time series data from  $X \in \mathbb{R}^{N \times T \times F}$  to  $X' \in \mathbb{R}^{N \times (T \times F)}$
  - 5: **Step 2: Calculate Global Feature Importance (GFI)**
  - 6: Train Random Forest regression on  $X_{\text{train}}$  to obtain feature importances  $GFI_i$  for each feature  $i$
  - 7: **Step 3: Select Neighbors for Anomaly Points**
  - 8: **for** each anomaly point  $x_a \in X_{\text{test}}$  **do**
  - 9:   **for** each candidate point  $x_c \in X_{\text{train}}$  **do**
  - 10:     Compute weighted cosine similarity:
  - 11:     
$$S(x_c, x_a) = \frac{\sum_i GFI'_i x_{ci} x_{ai}}{\sqrt{\sum_i GFI'_i x_{ci}^2} \sqrt{\sum_i GFI'_i x_{ai}^2}}$$
  - 12:     where weights  $GFI'_i = \exp(GFI_i)$
  - 13:   **end for**
  - 14:   Select top  $K$  candidate points with highest  $S(x_c, x_a)$
  - 15:   Set  $B_{x_a} = \{x_c \mid x_c \text{ is among top } K \text{ candidates}\}$
  - 16: **end for**
- 

#### 4.4. SHAP integration for anomalies

After the SHAP values are calculated, features need to be categorized into those that contribute to the anomaly (referred to as **contributors**) and those act as offsets. This categorization is based on the comparison between the actual observed output and the predicted value for a specific instance. Such categorization based on SHAP values helps us understand how much each feature pushes the model's prediction away from or towards the real value. The prediction for a sample is expressed as [21]:

$$\text{Prediction} = \text{BaseValue} + \sum_{i=1}^F \phi_i \quad (9)$$

Here, the term **BaseValue** represents the average prediction made using the background dataset, and the SHAP value  $\phi_i$  explains how much feature  $i$  contributes to the prediction. It is important to note that SHAP values are not used during the training of the model but are computed afterward to explain how features influence the final prediction.

The features are categorized according to the following logic:

- **If Real > Predicted:** This indicates that the model under-predicted the outcome. In this case:

$$\text{Real} > \text{BaseValue} + \sum_{i=1}^F \phi_i \quad (10)$$

Here, positive SHAP values ( $\phi_i > 0$ ) indicate features that pull the prediction closer to the real value, thereby acting as **offsets**. Negative SHAP values ( $\phi_i < 0$ ) indicate features that push the prediction further away from the true value, and these are considered **contributors** to the anomaly.

- **If Predicted > Real:** This indicates that the model over-predicted the outcome. In this case:

$$\text{BaseValue} + \sum_{i=1}^F \phi_i > \text{Real} \quad (11)$$

Positive SHAP values ( $\phi_i > 0$ ) represent features that push the prediction further from the real value, contributing to the over-prediction (**contributors**). Conversely, negative SHAP values ( $\phi_i < 0$ ) pull the prediction closer to the real value, acting as **offsets**.

In the previous steps, we reduced the contribution of many less effective features to near-zero values. However, some features still require further filtering, as shown in Fig. 7. By applying this categorization, we focus on features with negative SHAP values, which contribute to the anomaly. In this case, since the prediction is lower than the true value, the negative features (blue) are the main contributors, allowing us to disregard the positive features (red) for a more streamlined analysis. This process of classifying SHAP values is a part of SHAP integration for anomalies process as illustrated in Fig. 1.

## 5. Results and discussion

This section presents the outcomes of our explainable anomaly detection approach, including data preparation, model optimization, prediction performance, and feature importance. The results highlight the improvements achieved through hyperparameter tuning, the impact of key features on model predictions, and the advantages of using our approach for enhancing model explainability. Statistical comparisons further support the robustness of our approach, demonstrating its effectiveness in producing reliable and consistent results.

### 5.1. Dataset description and preparation

To conduct a comprehensive analysis of the proposed method, we use energy consumption datasets from five different consumer types: a residence, a manufacturing facility, a medical clinic, a retail store, and an office building. A residence dataset provided by London Hydro [49], comprises energy consumption records from a residence in London, Ontario, Canada, spanning from January 1, 2002, to December 31, 2004, with hourly energy consumption values. To enhance the predictive capabilities of anomaly detection and improve explanations as discussed in Subsection 4.1.1, we incorporated additional date-time and weather-related features from the Government of Canada's historical climate data [50]. The remaining four datasets are from Building Data Genome Project 2 [51], covering January 2016 to December 2017.



**Table 2**  
Summary of best hyperparameters and performance metrics for various models for the residential dataset.

Model	Best Parameters	MSE	RMSE	MAE	SMAPE	MAPE	R <sup>2</sup>
LSTM	88 LSTM units, dropout rate 0.224, 2 LSTM layers, learning rate 0.000279	0.09	0.30	0.15	9.23	8.98	0.41
GRU	88 GRU units, dropout rate 0.257, 1 GRU layer, learning rate 0.000997	<b>0.07</b>	<b>0.26</b>	<b>0.14</b>	<b>9.04</b>	<b>8.84</b>	<b>0.53</b>
BLSTM	109 LSTM units, dropout rate 0.285, 2 LSTM layers, learning rate 0.000189	0.09	0.30	0.15	9.18	8.96	0.43
BGRU	76 GRU units, dropout rate 0.272, 1 GRU layer, learning rate 0.000907	0.10	0.32	0.22	14.37	15.59	0.33
CNN	57 filters, kernel size 4, dropout rate 0.348, Adam optimizer, L2 regularization 1.01e-06	0.08	0.28	0.16	10.20	10.35	0.45
TCN	69 filters, kernel size 5, dropout rate 0.400, 5 TCN blocks, dilation base 3, L2 regularization 1.23e-06, Adam optimizer	0.08	0.28	0.17	10.95	11.41	0.46
DCNN	42 filters, kernel size 3, dilation rate 4, dropout rate 0.115, 1 convolutional layer, RMSprop optimizer, L2 regularization 7.37e-05	0.08	0.28	0.16	10.33	10.61	0.48
WaveNet	61 filters, kernel size 2, dilation rate 4, 2 WaveNet blocks, dropout rate 0.386, L2 regularization 7.84e-04, Adam optimizer	<b>0.07</b>	<b>0.26</b>	<b>0.14</b>	<b>9.04</b>	9.12	<b>0.53</b>
TFT	255 hidden units, dropout rate 0.312, 1 LSTM layer, 3 attention heads, 3 attention blocks, L2 regularization 8.81e-06, learning rate 0.000144, Adam optimizer	0.12	0.35	0.19	12.13	12.18	0.22
TST	Model dimension 65, 4 attention heads, 4 transformer layers, feed-forward dimension 471, dropout rate 0.306, learning rate 5.54e-05, Adam optimizer	0.11	0.33	0.22	13.79	14.69	0.29

Weather-related information for these datasets was already included in the repository.


The dataset was divided into training, validation, and test sets, following an 80-10-10 split. The model was trained based on windows sequences data. This configuration captures temporal dependencies and patterns within the data, significantly contributing to accurate anomaly detection in the analysis.

## 5.2. Optimization and training results

This study employs various deep learning architectures, including different variants of Recurrent Neural Networks (RNNs), Convolutional Neural Networks (CNNs), and Transformer-based models. The RNN variants utilized are Long Short-Term Memory (LSTM), Gated Recurrent Unit (GRU), and their bidirectional counterparts (BLSTM and BGRU), which effectively capture temporal dependencies in sequential data. Within the CNN category, we used one-dimensional CNN (1D-CNN), Dilated CNN (DCNN), Temporal Convolutional Networks (TCN), and WaveNet, which efficiently capture local temporal patterns and enhance feature extraction through convolutional operations. Additionally, Transformer-based models such as the Temporal Fusion Transformer (TFT) and Time Series Transformer (TST) were employed, leveraging self-attention mechanisms to dynamically weigh the importance of time steps and features.

Hyperparameter optimization was conducted for all algorithms using the Tree-structured Parzen Estimator (TPE) from Bayesian optimization. A total of 50 trials were performed, with each trial representing a different set of hyperparameters evaluated on the validation set. The best combination was selected to minimize validation loss. A summary of the selected hyperparameters for models trained on the residential dataset and performance metrics is presented in Table 2. For the remaining datasets, the same process was followed, but selected hyperparameters are not included for conciseness and Table 3 only includes performance metrics.

Next, the models with the best hyperparameter combinations were selected, and their performance was evaluated on the test set. Evaluation metrics included Mean Squared Error (MSE), Root Mean Squared Error (RMSE), Mean Absolute Error (MAE), Symmetric Mean Absolute Percentage Error (SMAPE), Mean Absolute Percentage Error (MAPE), and R<sup>2</sup>. Table 2 highlights the reliability of the prediction models on the residential dataset, which is essential for generating trustworthy explanations. Inaccurate models lead to unreliable explanations. The table



**Fig. 4.** Prediction results for the LSTM model, showing 10% of the test residential dataset. (For interpretation of the colors in the figure(s), the reader is referred to the web version of this article.)

presents various deep learning models with optimized hyperparameters, ensuring a fair performance comparison. Most models achieved high accuracy with low error rates, particularly the GRU model, which demonstrated the lowest errors and highest R<sup>2</sup> score. Fig. 4 shows an example of the LSTM model predictions compared to the actual values on 10% of the test data from the residential dataset, showing alignment between predicted and actual energy consumption.

To examine the performance of the anomaly detection models on the remaining four datasets, a manufacturing facility, a medical clinic, a retail store, and an office building, Table 3 presents the key performance metrics for each model type. Note that metrics such as MSE, RMSE, and MAE should not be compared among datasets and those are scale-dependent metrics. No single algorithm outperforms all others across all performance metrics and datasets; nevertheless, TCN achieves overall excellent performance. We evaluate our explainability approach across all listed models.

## 5.3. Global feature importance

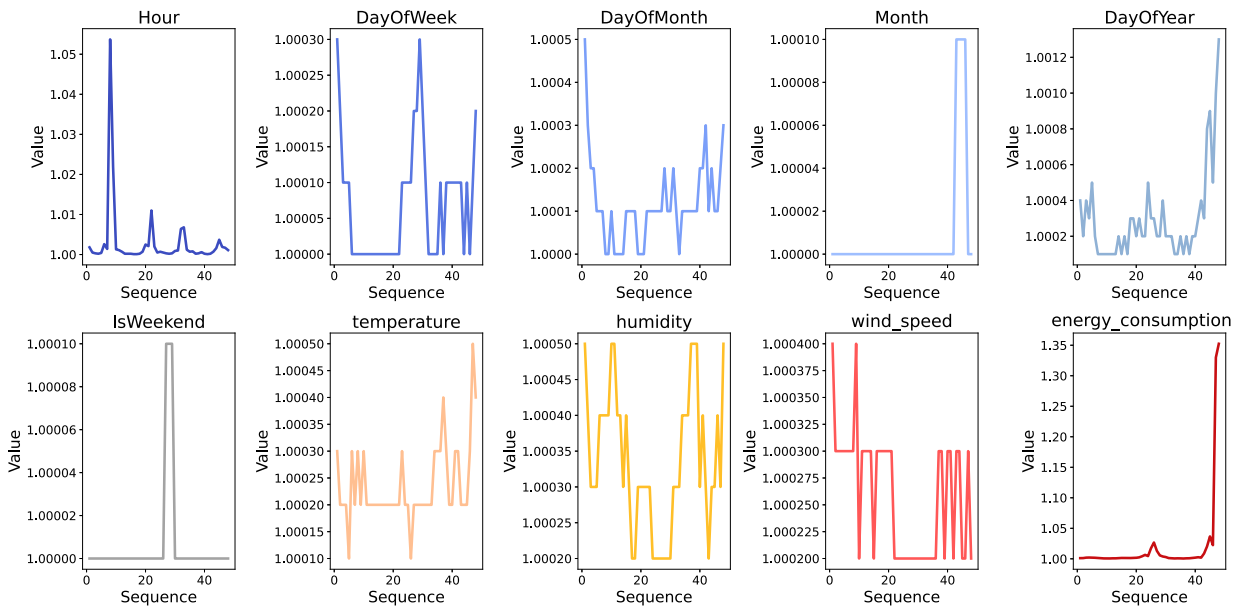
In this section, we present the global feature importance results for 10 features across 48 time sequences, offering a comprehensive view of each feature's contribution to model prediction. As detailed in the methodology, we applied an exponential transformation to the Random Forest results to better highlight the distinctions in feature importance. Based on Fig. 5 for the residential data, the previous hours energy con-

**Table 3**  
Summary of performance metrics for additional datasets.

(a) Manufacturing Facility							(b) Medical Clinic						
Model	MSE	RMSE	MAE	SMAPE	MAPE	R <sup>2</sup>	Model	MSE	RMSE	MAE	SMAPE	MAPE	R <sup>2</sup>
LSTM	254.49	15.95	9.18	11.21	11.42	0.9	LSTM	9.56	3.09	2.28	12.04	11.95	0.87
GRU	306.9	17.52	11.03	14.33	14.5	0.88	GRU	9.04	3.01	2.25	11.83	12.08	0.88
BLSTM	291.82	17.08	9.93	12.13	12.63	0.88	BLSTM	8.1	2.85	2.15	11.34	11.5	0.89
BGRU	318.79	17.85	10.6	12.96	13.08	0.87	BGRU	8.64	2.94	2.14	11.15	11.2	0.89
CNN	201.35	14.19	7.85	9.89	10.15	<b>0.92</b>	CNN	12.09	3.48	2.51	12.36	11.65	0.84
TCN	200.1	14.15	8.16	10.31	10.9	<b>0.92</b>	TCN	<b>5.19</b>	<b>2.28</b>	<b>1.62</b>	<b>8.48</b>	<b>8.46</b>	<b>0.93</b>
DCNN	209.39	14.47	8.07	9.88	10.27	<b>0.92</b>	DCNN	7.23	2.69	1.97	10.19	9.87	0.91
WaveNet	<b>188.73</b>	<b>13.74</b>	<b>7.8</b>	<b>9.88</b>	<b>10.15</b>	<b>0.92</b>	WaveNet	5.26	2.29	1.63	8.75	8.76	<b>0.93</b>
TFT	449.58	21.2	10.98	12.61	13.01	0.82	TFT	23.67	4.87	3.39	17.39	16.63	0.69
TST	1031.83	32.12	19	20.44	21.5	0.58	TST	9.63	3.1	2.3	11.98	12.36	0.87

(c) Retail Store							(d) Office Building						
Model	MSE	RMSE	MAE	SMAPE	MAPE	R <sup>2</sup>	Model	MSE	RMSE	MAE	SMAPE	MAPE	R <sup>2</sup>
LSTM	0.49	0.7	0.45	15.45	16.45	0.84	LSTM	44.42	6.66	4.81	11.5	12.11	0.86
GRU	0.53	0.73	0.5	17.33	18.76	0.83	GRU	40.8	6.39	4.65	11.17	11.89	0.87
BLSTM	0.55	0.74	0.47	15.8	16.46	0.82	BLSTM	48.7	6.98	5.13	12.68	13.75	0.85
BGRU	0.49	0.7	0.47	16.41	16.97	0.84	BGRU	38.19	6.18	4.5	10.94	11.36	0.88
CNN	0.42	0.65	0.44	15.51	16.85	<b>0.86</b>	CNN	23.33	4.83	3.59	8.69	8.94	0.93
TCN	0.43	0.66	<b>0.43</b>	15.15	<b>15.18</b>	<b>0.86</b>	TCN	20.99	4.58	3.35	8.12	8.26	0.93
DCNN	<b>0.41</b>	<b>0.64</b>	<b>0.43</b>	15.44	16.24	<b>0.86</b>	DCNN	26.19	5.12	3.8	9.4	9.49	0.92
WaveNet	0.42	0.65	<b>0.43</b>	15.6	15.96	<b>0.86</b>	WaveNet	<b>19.84</b>	<b>4.45</b>	<b>3.22</b>	<b>7.88</b>	<b>8.05</b>	<b>0.94</b>
TFT	1.04	1.02	0.66	21.7	24.46	0.66	TFT	72.87	8.54	6.16	14.2	14.83	0.77
TST	0.47	0.69	0.44	<b>14.89</b>	15.96	0.84	TST	95.46	9.77	6.97	15.62	17	0.7



**Fig. 5.** Global Feature Importance Line Plot of Sequence Windows for LSTM Model on residential dataset.

sumption data emerged as the most significant predictor for the LSTM model. The energy consumption feature exhibited a consistently increasing trend in importance over the last 48 hours, peaking at values of 1.352 and 1.329 for the final time steps, highlighting its relevance near the prediction window (e.g., sequence of 48 hours before prediction). The hour feature also exhibited notable influence, with peak values of 1.0068 and 1.0064. Other features, while less influential, follow consistent temporal trends. These observations highlight the importance of key features in selecting a contextually relevant baseline for explainability. This process was repeated for other deep learning models and other datasets, consistently revealing that energy consumption history and time-related features are critical for predictions due to the inherent temporal patterns in energy consumption. These global feature importance results will guide the following sections, where we focus on selecting a background dataset that aligns with the most important features identified

in the global model. This selection will enhance the explainability process by concentrating on the features that play a key role in explaining anomalies.

**5.4. Explanation step: comparing random background dataset and similar background dataset selection**

In this section, we provide a detailed comparison between the random and similar background dataset selection methods by analyzing a specific anomaly point identified using the LSTM model on the residential data, as shown in Figs. 6 and 7.

As discussed in Section 5.3, these feature importance scores have dimensions of 48×10, where the features are represented in the rows, and their sequences are displayed along the columns. The features are sorted by the absolute value of their SHAP values, with the most impactful fea-




Fig. 6. SHAP value heatmap for random background dataset selection approach for LSTM Model on residential dataset.




Fig. 7. SHAP value heatmap for similar background dataset selection approach for LSTM Model on residential dataset.

tures positioned at the top. For example, in Fig. 6, the date and time features, such as the month, are identified as the most important in explaining the anomaly. In contrast, in Fig. 7, energy consumption and weather-related features are more significant contributors. This demonstrates that the similar background dataset selection can capture different aspects of the data, such as weather information, due to its temporal alignment.

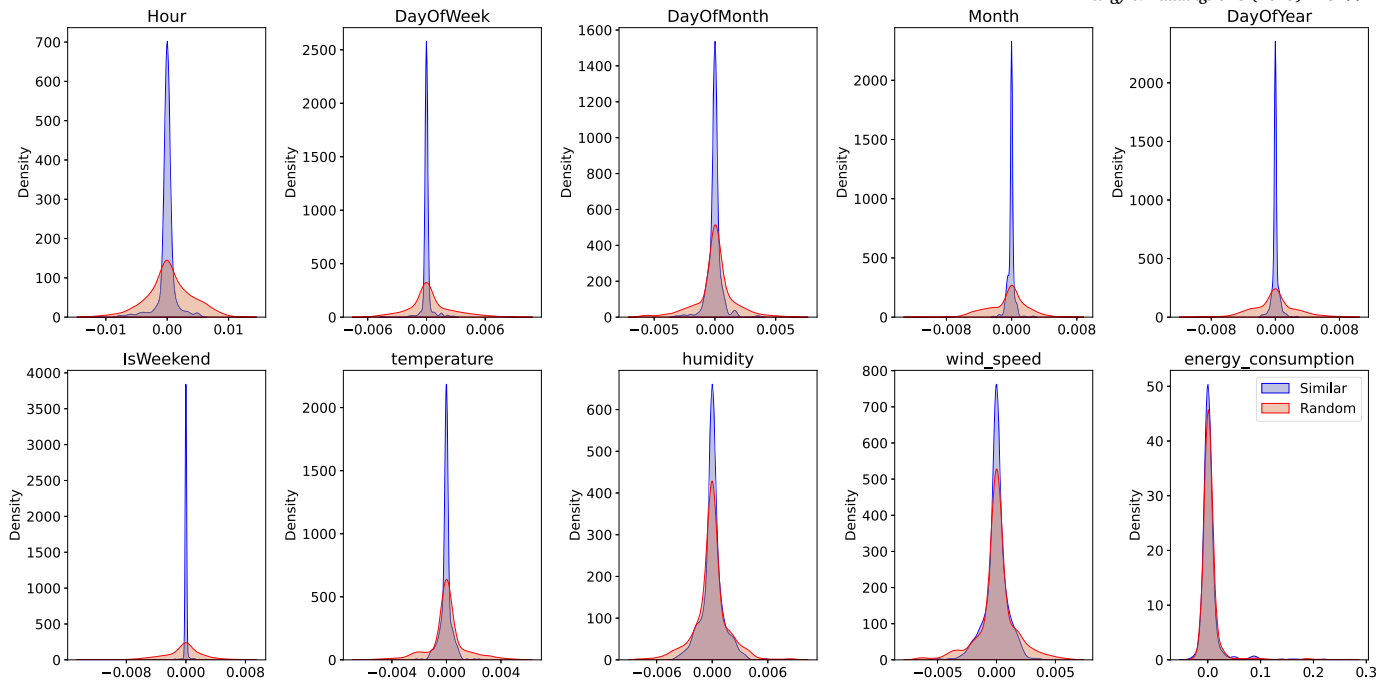
Another important aspect of these heatmaps is their relevance to the reliability of the results. Before explaining this further, it is important to clarify the role of the line plots and the dotted horizontal line present at the top of each heatmap. As outlined in Equation 9, the prediction function  $f(x)$  (represented by the line plot) is the sum of the SHAP values, showing how each feature shifts the prediction either away from or towards the base value (dotted horizontal line). Given that the dataset consists of 10 features and 48 time steps, the prediction function is defined as:

$$f(x) = \text{BaseValue} + \sum_{t=1}^{48} \sum_{i=1}^{10} \text{SHAP}_t(i) \quad (12)$$

Here, the *Base Value* represents the average prediction derived from the background dataset, and  $\text{SHAP}_t(i)$  represents the contribution of feature  $i$  at time step  $t$ . This breakdown provides a clear view of how each feature at each time step influences the overall prediction.

From the heatmaps in Figs. 6 and 7, which correspond to the same anomaly detected by the LSTM model, the prediction, true value, and error are 1.60, 4.75, and -3.15, respectively. Since the model's prediction is lower than the true value, positive SHAP values (red cells) represent features that adjust the prediction closer to the true value (offsets), while negative SHAP values (blue cells) correspond to features that push the prediction further from the true value, contributing to the anomaly.

In the following, we compare the random and similar background dataset selection based on these aspects:



**Fig. 8.** Density plots of SHAP values for similar vs. random background dataset across 10 features, based on anomalies detected by the LSTM model on residential dataset.

- **Random Background Dataset (Fig. 6):** This heatmap shows significant variability in the SHAP values across the sequence of features, with many features showing large positive (bright red) or negative (bright blue) impacts. This variability causes fluctuations in the prediction function  $f(x)$  around the base value (dotted line), making it difficult to clearly identify the most important features. The high variability suggests that the randomly selected background dataset may not align closely with the anomaly point, potentially overestimating the importance of less relevant features, such as date and time attributes.
- **Similar Background Dataset (Fig. 7):** In contrast, the heatmap from the similar background dataset shows a more stable pattern (fewer fluctuations of  $f(x)$  around the base value), with energy consumption and weather-related features emerging as the primary drivers of the model's predictions. This approach reduces the noise caused by less relevant features (more features near zero as most of them with light red or blue colors), offering a clearer identification of the factors contributing to the anomaly. For instance, energy consumption functions as an offset, correcting the model's prediction towards the true value, while weather features such as humidity, temperature, and wind speed reflect genuine variations that contribute to the anomaly.

These observations underscore the effectiveness of using similar background dataset selected according to the proposed similarity metrics in enhancing the interpretability of SHAP values. By focusing on a relevant background dataset, we can better isolate the features that genuinely influence the model's predictions, reducing the impact of unrelated features. This approach is especially valuable given the complexity of the dataset, which includes 480 features (48 sequences  $\times$  10 features). Using appropriate background dataset allows for more focused and meaningful explanations of the model's behavior.

### 5.5. SHAP density plot analysis

As seen from the density plots in Fig. 8, based on anomalies detected by the LSTM model on residential dataset, a notable reduction in the variation of SHAP values is observed for date and time-related features

(Hour, DayOfWeek, DayOfMonth, Month, DayOfYear) when utilizing similar background dataset. This reduction underscores the ability of this approach to filter out less relevant features, thereby allowing the model to focus on factors more closely linked to the detected anomalies, such as weather information. However, while the variation is significantly reduced, these features are not entirely disregarded. In cases where energy consumption is unusual at specific times, these features still contribute to the model's prediction, albeit with less fluctuation.

This behavior indicates that the similar background dataset essentially serves as a filter, helping to emphasize the most impactful features while minimizing noise from less relevant ones. This filtering effect is particularly beneficial when working with high-dimensional data, where many features may not be directly related to the anomalies. By concentrating the SHAP values near zero for these date and time features, the model can more effectively highlight the key drivers of unusual energy consumption, such as weather conditions or specific patterns in energy use.

### 5.6. Stability of SHAP explanations across multiple datasets and explainability approaches

Analysis of results so far focused on a residential dataset, demonstrating the benefit of our technique for achieving robust explanations. Building on this, here we expand the analysis to the remaining four datasets— a manufacturing facility, a medical clinic, a retail store, and an office building—to examine the portability of our technique across diverse energy consumer types. Additionally, for each dataset, we consider diverse XAI algorithms to illustrate the wide applicability of our technique.

Moreover, to further evaluate the impact of our approach for background dataset selection methods on SHAP values, we conducted a comparison between random and similar background dataset approaches across different deep learning models and datasets from diverse energy consumer types. This analysis explored various explainability algorithms, including SHAP variants (Kernel, Partition, and Sampling), as well as LIME and Permutation, applied across all five datasets.

Table 4 presents the results. For each dataset and each XAI method, the table reports the mean and standard deviation (Mean  $\pm$  SD) of

**Table 4**  
Comparison of variance reduction and statistical results across datasets and XAI methods.

Dataset	XAI methods	Random Baseline (mean $\pm$ sd)	Proposed Baseline (mean $\pm$ sd)	Statistical Test (P-Value)	Reduction (%)
Residential	Kernel	0.050 $\pm$ 0.110	0.028 $\pm$ 0.034	27.851 (0.009*)	44.0
	Lime	0.077 $\pm$ 0.192	0.040 $\pm$ 0.075	27.773 (0.119)	48.1
	Partition	0.023 $\pm$ 0.064	0.017 $\pm$ 0.026	219.271 (0.138)	26.1
	Permutation	0.045 $\pm$ 0.105	0.024 $\pm$ 0.031	27.326 (0.029*)	46.7
	Sampling	0.045 $\pm$ 0.105	0.024 $\pm$ 0.031	27.646 (0.001**)	46.7
Manufacturing facility	Kernel	1.214 $\pm$ 1.932	0.360 $\pm$ 0.667	23.598 (0.000**)	70.3
	Lime	0.478 $\pm$ 1.233	0.377 $\pm$ 1.015	8.857 (0.243)	21.1
	Partition	0.238 $\pm$ 0.766	0.047 $\pm$ 0.113	59.502 (0.000**)	80.3
	Permutation	0.759 $\pm$ 2.046	0.328 $\pm$ 0.681	25.524 (0.000**)	56.8
	Sampling	0.739 $\pm$ 2.037	0.313 $\pm$ 0.667	24.407 (0.000**)	57.6
Medical clinic	Kernel	0.126 $\pm$ 0.267	0.081 $\pm$ 0.125	17.410 (0.028*)	35.7
	Lime	0.117 $\pm$ 0.290	0.096 $\pm$ 0.226	2.562 (0.242)	17.9
	Partition	0.043 $\pm$ 0.134	0.035 $\pm$ 0.084	12.215 (0.198)	18.6
	Permutation	0.116 $\pm$ 0.264	0.074 $\pm$ 0.121	17.969 (0.039*)	36.2
	Sampling	0.115 $\pm$ 0.262	0.073 $\pm$ 0.121	17.622 (0.015*)	36.5
Retail store	Kernel	0.034 $\pm$ 0.046	0.023 $\pm$ 0.036	2.662 (0.567)	32.4
	Lime	0.024 $\pm$ 0.037	0.023 $\pm$ 0.030	2.391 (0.453)	4.2
	Partition	0.015 $\pm$ 0.014	0.010 $\pm$ 0.000	757.685 (0.000**)	33.3
	Permutation	0.031 $\pm$ 0.049	0.024 $\pm$ 0.036	2.938 (0.530)	22.6
	Sampling	0.031 $\pm$ 0.047	0.023 $\pm$ 0.036	2.856 (0.563)	25.8
Office building	Kernel	0.219 $\pm$ 0.410	0.105 $\pm$ 0.182	24.249 (0.002**)	52.1
	Lime	0.244 $\pm$ 0.577	0.156 $\pm$ 0.306	15.670 (0.038*)	36.1
	Partition	0.019 $\pm$ 0.033	0.017 $\pm$ 0.024	38.760 (0.133)	10.5
	Permutation	0.195 $\pm$ 0.426	0.104 $\pm$ 0.183	26.349 (0.000**)	46.7
	Sampling	0.201 $\pm$ 0.448	0.105 $\pm$ 0.184	28.660 (0.000**)	47.8

\*\* indicates significance at the 1% level; \* indicates significance at the 5% level.

SHAP/LIME value variability for random baseline and our baseline, calculated as the standard deviation of feature importance across detected anomalies. For all datasets and all XAI methods, our technique achieves lower variability than the random baseline, highlighting its ability to provide more stable and reliable explanations compared to the random baseline. Additionally, the table includes the percentage reduction in variability, calculated as the relative decrease in the similar baseline's mean compared to the random baseline, to quantify the effectiveness of the proposed method. Bartlett's test results (P-Value) [52] are also presented to assess the statistical significance of the differences in variability between the two approaches, random baseline and our baseline. A single \* indicates significance on the 5% level and \*\* denotes significance on the 1% level.

For the residential dataset, significant reductions in variability were observed with Kernel (44%,  $P = 0.009$ ), Permutation (46.7%,  $P = 0.029$ ), and Sampling SHAP (46.7%,  $P = 0.001$ ), while LIME (48.1%,  $P = 0.119$ ) and Partition (26.1%,  $P = 0.138$ ) showed notable reductions without statistical significance. In the manufacturing facility dataset, the proposed method achieved particularly strong results, with Partition (80.3%,  $P < 0.001$ ), Kernel (70.3%,  $P < 0.001$ ), Permutation (56.8%,  $P < 0.001$ ), and Sampling SHAP (57.6%,  $P < 0.001$ ) all demonstrating significant variability reductions. LIME, although less impactful, still achieved a reduction of 21.1% ( $P = 0.243$ ).

The medical clinic dataset showed significant improvements for Kernel (35.7%,  $P = 0.028$ ), Permutation (36.2%,  $P = 0.039$ ), and Sampling SHAP (36.5%,  $P = 0.015$ ), with LIME (17.9%,  $P = 0.242$ ) and Partition SHAP (18.6%,  $P = 0.198$ ) exhibiting modest reductions. The retail store dataset presented mixed results, with statistical significance only achieved for Partition SHAP (33.3%,  $P < 0.001$ ). Kernel (32.4%,  $P = 0.567$ ), Sampling (25.8%,  $P = 0.563$ ), and Permutation SHAP (22.6%,  $P = 0.530$ ) demonstrated modest reductions without significance, while LIME showed minimal impact (4.2%,  $P = 0.453$ ). Finally, for the office building dataset, significant improvements were observed with Kernel (52.1%,  $P = 0.002$ ), Permutation (46.7%,  $P < 0.001$ ), and Sampling SHAP (47.8%,  $P < 0.001$ ), while LIME (36.1%,  $P = 0.038$ ) also achieved

a statistically significant reduction. Partition SHAP, however, showed limited impact with a reduction of 10.5% ( $P = 0.133$ ).

Overall, our approach reduced the variability of the explanations across all datasets and across all XAI methods, with most differences being statistically significant. Particularly good improvements were observed for Kernel, Permutation, and Sampling SHAP. For some datasets and XAI methods, the reduction was larger than for others, but it was present for all datasets. These findings underscore the robustness of the baseline selection strategy in reducing variability and enhancing the stability of SHAP explanations across multiple datasets and XAI methods.

## 6. Conclusion

This study presents a comprehensive methodology for explaining deep learning-based anomaly detection in energy consumption data. Through a systematic approach that integrates advanced anomaly detection techniques and innovative explanation methods leveraging a novel approach for selecting the background dataset (baseline) for model-agnostic explainability algorithms, the study provides improved explainability of detected anomalies applicable across a diverse range of deep learning models, SHAP variants, and energy consumption datasets. The key innovation lies in the selection of context-relevant information for the baseline employed in explaining anomalies. Moreover, we guide the explanations toward features highly relevant in the deep learning model. This strategy ensures more consistent explanations, enabling a deeper understanding of the factors driving anomaly detection.

The evaluation conducted on five datasets, with five XAI approaches (Kernel SHAP, Partition SHAP, Sampling SHAP, LIME, Permutation) demonstrated that our baseline approach significantly reduces variability in feature attributions, with reductions ranging from 26.1% to 80.3%, depending on the dataset and a SHAP variant. The statistical analysis validated that our approach, compared to the random baseline, achieves a significant reduction in variability for Kernel, Permutation, and Sampling SHAP across most datasets. Although LIME and Partition SHAP did not always achieve statistical significance, they still exhibited reductions in variability.

Future work will examine developing additional metrics for evaluating the quality of explanations as well as examining human perception of the evaluations. Moreover, the approach will be examined with different use cases.

### CRedit authorship contribution statement

**Mohammad Noorchenarboo:** Writing – review & editing, Writing – original draft, Software, Methodology, Data curation, Conceptualization. **Katarina Grolinger:** Writing – review & editing, Supervision, Methodology, Funding acquisition, Conceptualization.

### Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: Katarina Grolinger reports financial support was provided by Natural Sciences and Engineering Research Council of Canada. If there are other authors, they declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

### Acknowledgement

This work was supported by London Hydro and Natural Sciences and Engineering Research Council of Canada (NSERC) under grant ALLRP 571890 - 21. Computation was enabled in part by the Digital Research Alliance of Canada.

### Data availability

The data that has been used is confidential.

### References

- [1] C. Fan, F. Xiao, Y. Zhao, J. Wang, Analytical investigation of autoencoder-based methods for unsupervised anomaly detection in building energy data, *Appl. Energy* 211 (2018) 1123–1135, <https://doi.org/10.1016/j.apenergy.2017.12.005>.
- [2] I. Martin-Hoyo, F. Bouly, N. Gandolfo, C. Joly, R. Paparella, J.-L. Biarrotte, J.M. Martin-Sanchez, Optimized adaptive control for the MYRRA linear accelerator: control system design for a superconducting cavity in a particle accelerator, *IEEE Control Syst. Mag.* 38 (2) (2018) 44–79, <https://doi.org/10.1109/mcs.2017.2786420>.
- [3] H. Pan, Z. Yin, X. Jiang, High-dimensional energy consumption anomaly detection: a deep learning-based method for detecting anomalies, *Energies* 15 (17) (2022) 6139, <https://doi.org/10.3390/en15176139>.
- [4] W. Dai, X. Liu, A. Heller, P.S. Nielsen, Smart meter data anomaly detection using variational recurrent autoencoders with attention, in: *Intelligent Technologies and Applications*, Springer International Publishing, 2022, pp. 311–324.
- [5] X. Liu, Y. Ding, H. Tang, F. Xiao, A data mining-based framework for the identification of daily electricity usage patterns and anomaly detection in building electricity consumption data, *Energy Build.* 231 (2021) 110601, <https://doi.org/10.1016/j.enbuild.2020.110601>.
- [6] Y. Himeur, K. Ghanem, A. Alsalemi, F. Bensaali, A. Amira, Artificial intelligence based anomaly detection of energy consumption in buildings: a review, current trends and new perspectives, *Appl. Energy* 287 (2021) 116601, <https://doi.org/10.1016/j.apenergy.2021.116601>.
- [7] C. Xu, H. Chen, A hybrid data mining approach for anomaly detection and evaluation in residential buildings energy data, *Energy Build.* 215 (2020) 109864, <https://doi.org/10.1016/j.enbuild.2020.109864>.
- [8] F. Dehrouyeh, L. Yang, F. Badrkhani Ajaei, A. Shami, On TinyML and cybersecurity: electric vehicle charging infrastructure use case, *IEEE Access* 12 (2024) 108703–108730, <https://doi.org/10.1109/access.2024.3437192>.
- [9] G. Li, J.J. Jung, Deep learning for anomaly detection in multivariate time series: approaches, applications, and challenges, *Inf. Fusion* 91 (2023) 93–102, <https://doi.org/10.1016/j.inffus.2022.10.008>.
- [10] H. Esen, F. Ozgen, M. Esen, A. Sengur, Artificial neural network and wavelet neural network approaches for modelling of a solar air heater, *Expert Syst. Appl.* 36 (8) (2009) 11240–11248, <https://doi.org/10.1016/j.eswa.2009.02.073>.
- [11] H. Esen, M. Esen, O. Ozsolak, Modelling and experimental performance analysis of solar-assisted ground source heat pump system, *J. Exp. Theor. Artif. Intell.* 29 (1) (2017) 1–17, <https://doi.org/10.1080/0952813x.2015.1056242>.
- [12] Z. Qu, H. Liu, Z. Wang, J. Xu, P. Zhang, H. Zeng, A combined genetic optimization with adaboost ensemble model for anomaly detection in buildings electricity consumption, *Energy Build.* 248 (2021) 111193, <https://doi.org/10.1016/j.enbuild.2021.111193>.
- [13] H. Esen, M. Inalli, A. Sengur, M. Esen, Artificial neural networks and adaptive neuro-fuzzy assessments for ground-coupled heat pump system, *Energy Build.* 40 (6) (2008) 1074–1083, <https://doi.org/10.1016/j.enbuild.2007.10.002>.
- [14] V. Yepmo, G. Smits, O. Pivert, Anomaly explanation: a review, *Data Knowl. Eng.* 137 (2022) 101946, <https://doi.org/10.1016/j.datak.2021.101946>.
- [15] L. Antwarg, R.M. Miller, B. Shapira, L. Rokach, Explaining anomalies detected by autoencoders using Shapley additive explanations, *Expert Syst. Appl.* 186 (2021) 115736, <https://doi.org/10.1016/j.eswa.2021.115736>.
- [16] A. Ravi, X. Yu, I. Santelices, F. Karray, B. Fidan, General frameworks for anomaly detection explainability: comparative study, in: *2021 IEEE International Conference on Autonomous Systems (ICAS)*, vol. 30, IEEE, 2021, pp. 1–5.
- [17] U. Schlegel, H. Arnout, M. El-Assady, D. Oelke, D.A. Keim, Towards a rigorous evaluation of XAI methods on time series, in: *2019 IEEE/CVF International Conference on Computer Vision Workshop (ICCVW)*, IEEE, 2019, pp. 4197–4201.
- [18] S.M. Lundberg, S.-I. Lee, A unified approach to interpreting model predictions, in: *31st International Conference on Neural Information Processing Systems*, Curran Associates Inc., 2017, pp. 4768–4777, <https://arxiv.org/abs/1705.07874>.
- [19] D. Fryer, I. Strumke, H. Nguyen, Shapley values for feature selection: the good, the bad, and the axioms, *IEEE Access* 9 (2021) 144352–144360, <https://doi.org/10.1109/access.2021.3119110>.
- [20] G. Kelodjou, L. Rozé, V. Masson, L. Galárraga, R. Gaudel, M. Tchuente, A. Termier, Shaping up SHAP: enhancing stability through layer-wise neighbor selection, in: *AAAI Conference on Artificial Intelligence*, vol. 38, Association for the Advancement of Artificial Intelligence (AAAI), 2024, pp. 13094–13103.
- [21] K. Aas, M. Jullum, A. Leland, Explaining individual predictions when features are dependent: more accurate approximations to Shapley value, *Artif. Intell.* 298 (2021) 103502, <https://doi.org/10.1016/j.artint.2021.103502>.
- [22] H. Chen, I.C. Covert, S.M. Lundberg, S.-I. Lee, Algorithms to estimate Shapley value feature attributions, *Nat. Mach. Intell.* 5 (6) (2023) 590–601, <https://doi.org/10.1038/s42256-023-00657-x>.
- [23] J. Moon, S. Rho, S.W. Baik, Toward explainable electrical load forecasting of buildings: a comparative study of tree-based ensemble methods with Shapley values, *Sust. Energy Technol. Assess.* 54 (2022) 102888, <https://doi.org/10.1016/j.seta.2022.102888>.
- [24] W.J. Chung, C. Liu, Analysis of input parameters for deep learning-based load prediction for office buildings in different climate zones using explainable artificial intelligence, *Energy Build.* 276 (2022) 112521, <https://doi.org/10.1016/j.enbuild.2022.112521>.
- [25] K. Joshi, A. Jana, P. Arjunan, K. Ramaritham, XENIA: explainable energy informatics and attributes for building energy benchmarking, in: *9th ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation*, ACM, 2022, pp. 406–412.
- [26] K. Zhang, P. Xu, J. Zhang, Explainable AI in deep reinforcement learning models: a SHAP method applied in power system emergency control, in: *2020 IEEE 4th Conference on Energy Internet and Energy System Integration (EI2)*, IEEE, 2020, pp. 419–427.
- [27] B. Tan, J. Zhao, T. Su, Q. Huang, Y. Zhang, H. Zhang, Explainable Bayesian neural network for probabilistic transient stability analysis considering wind energy, in: *2022 IEEE Power & Energy Society General Meeting (PESGM)*, IEEE, 2022, pp. 1–5.
- [28] D. Leuthe, J. Mirlach, S. Wenninger, C. Wieth, Leveraging explainable AI for informed building retrofit decisions: insights from a survey, *Energy Build.* 318 (2024) 114426, <https://doi.org/10.1016/j.enbuild.2024.114426>.
- [29] J. Mueller, L. Czekalla, F. Schuchter, G. Frey, Illuminating the black box: a comparative study of explainable AI for interpreting time series data in vehicle power net consumption models, in: *2023 International Conference on Machine Learning and Applications (ICMLA)*, vol. 70, IEEE, 2023, pp. 166–173.
- [30] J. Rožanec, E. Trajkova, K. Kenda, B. Fortuna, D. Mladenčić, Explaining bad forecasts in global time series models, *Appl. Sci.* 11 (19) (2021) 9243, <https://doi.org/10.3390/app11199243>.
- [31] J. Labaien Soto, E. Zugasti Uriguén, X. De Carlos Garcia, Real-time, model-agnostic and user-driven counterfactual explanations using autoencoders, *Appl. Sci.* 13 (5) (2023) 2912, <https://doi.org/10.3390/app13052912>.
- [32] W. Freeborough, T. van Zyl, Investigating explainability methods in recurrent neural network architectures for financial time series data, *Appl. Sci.* 12 (3) (2022) 1427, <https://doi.org/10.3390/app12031427>.
- [33] U. Schlegel, D.A. Keim, A deep dive into perturbations as evaluation technique for time series XAI, in: *Explainable Artificial Intelligence*, Springer Nature Switzerland, 2023, pp. 165–180.
- [34] K. Roshan, A. Zafar, Utilizing XAI technique to improve autoencoder based model for computer network anomaly detection with Shapley additive explanation (SHAP), *Int. J. Comput. Netw. Commun.* 13 (6) (2021) 109–128, <https://doi.org/10.5121/ijnc.2021.13607>.
- [35] D. Kim, G. Antariksa, M.P. Handayani, S. Lee, J. Lee, Explainable anomaly detection framework for maritime main engine sensor data, *Sensors* 21 (15) (2021) 5200, <https://doi.org/10.3390/s21155200>.

- [36] M. Abououf, S. Singh, R. Mizouni, H. Otrok, Explainable AI for event and anomaly detection and classification in healthcare monitoring systems, *IEEE Int. Things J.* 11 (2) (2024) 3446–3457, <https://doi.org/10.1109/jiot.2023.3296809>.
- [37] S. Asutkar, S. Tallur, An explainable unsupervised learning framework for scalable machine fault detection in industry 4.0, *Meas. Sci. Technol.* 34 (10) (2023) 105123, <https://doi.org/10.1088/1361-6501/ace640>.
- [38] S.H. Choi, J.M. Lee, Explainable fault diagnosis model using stacked autoencoder and kernel SHAP, in: *2022 IEEE International Symposium on Advanced Control of Industrial Processes (AdCONIP)*, vol. 55812, IEEE, 2022, pp. 182–187.
- [39] B. Goodman, S. Flaxman, European Union regulations on algorithmic decision making and a “right to explanation”, *AI Mag.* 38 (3) (2017) 50–57, <https://doi.org/10.1609/aimag.v38i3.2741>.
- [40] D. Castelvechi, Can we open the black box of AI?, *Nature* 538 (7623) (2016) 20–23, <https://doi.org/10.1038/538020a>.
- [41] N. Spyrisson, D. Cook, P. Biecek, Exploring local explanations of nonlinear models using animated linear projections, *Comput. Stat.* (2024), <https://doi.org/10.1007/s00180-023-01453-2>.
- [42] G. Fernandez, R. Guidotti, F. Giannotti, M. Setzu, J.A. Aledo, J.A. Gámez, J.M. Puerta, Flocalx - local to global fuzzy explanations for black box classifiers, in: *Advances in Intelligent Data Analysis XXII*, Springer Nature Switzerland, 2024, pp. 197–209.
- [43] M.R. Zafar, N. Khan, Deterministic local interpretable model-agnostic explanations for stable explainability, *Mach. Learn. Knowl. Extr.* 3 (2021) 525–541, <https://doi.org/10.3390/make3030027>.
- [44] S. Lipovetsky, M. Conklin, Analysis of regression in game theory approach, *Appl. Stoch. Models Bus. Ind.* 17 (4) (2001) 319–330, <https://doi.org/10.1002/asmb.446>.
- [45] T. Akiba, S. Sano, T. Yanase, T. Ohta, M. Koyama, Optuna: a next-generation hyperparameter optimization framework, in: *25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, vol. 18, ACM, 2019, pp. 2623–2631.
- [46] D. Dallah, H. Sulieman, Outlier detection using the range distribution, in: *Advances in Mathematical Modeling and Scientific Computing*, Springer International Publishing, 2024, pp. 687–697.
- [47] S. Nembrini, I.R. König, M.N. Wright, The revival of the Gini importance?, *Bioinformatics* 34 (21) (2018) 3711–3718, <https://doi.org/10.1093/bioinformatics/bty373>.
- [48] H. Wang, F. Yang, Z. Luo, An experimental study of the intrinsic stability of random forest variable importance measures, *BMC Bioinform.* 17 (1) (2016), <https://doi.org/10.1186/s12859-016-0900-5>.
- [49] London hydro, <https://www.londonhydro.com/>. (Accessed 26 September 2024).
- [50] Historical climate data - climate - environment and climate change Canada, [https://climate.weather.gc.ca/historical\\_data/search\\_historic\\_data\\_e.html](https://climate.weather.gc.ca/historical_data/search_historic_data_e.html). (Accessed 26 September 2024).
- [51] C. Miller, A. Kathirgamanathan, B. Picchetti, P. Arjunan, J.Y. Park, Z. Nagy, P. Raftery, B.W. Hobson, Z. Shi, F. Meggers, The building data genome project 2, energy meter data from the ashrae great energy predictor III competition, *Sci. Data* 7 (1) (2020), <https://doi.org/10.1038/s41597-020-00712-x>.
- [52] M.S. Bartlett, *Properties of Sufficiency and Statistical Tests*, Springer, New York, 1992, pp. 113–126.