ORIGINAL RESEARCH **Open access**

Interpretable Attention-Based Autoencoder for Detecting Energy Anomalies in Buildings

Jan Peeters[1*], Marie Janssens[1], Wouter Claes[2]

Abstract

Detecting departures from standard operational states promptly, coupled with corrective interventions, can yield massive energy reductions in the built environment. To automate the recognition of these abnormal events, building anomaly detection frameworks exploit the complex, high-dimensional records acquired by smart utility meters. While autoencoders excel at compressing dimensions and mapping the fundamental data distribution, they inherently process individual features in isolation. To overcome this limitation, the present research integrates an attention mechanism within an autoencoder topology to account for cross-feature correlations explicitly. The advantages of this attention framework are demonstrated by benchmarking its reconstructive performance against a conventional ANN-based autoencoder using artificially generated data. The findings reveal that incorporating an attention layer grants the encoder–decoder system greater immunity to training data outliers, which in turn minimizes the need for data preprocessing. Additionally, the proposed architecture is validated on an empirical dataset. The resulting attention maps help decode the dynamic interactions across both the temporal and feature spaces, ultimately delivering a transparent, human-readable methodology for understanding the model's reconstructive outputs.

Keywords Artificial intelligence, UN Sustainable Development Goals, Interpretable model, Multivariate time series, HVAC

*Correspondence:
Jan Peeters
jan.peeters@kuleuven.be
[1] Department of Intelligent Energy Management Systems, Faculty of Engineering, KU Leuven, Leuven, Belgium
[2] Department of AI for Hybrid Renewable Systems, Faculty of Engineering, Ghent University, Ghent, Belgium

Introduction

Background

As detailed in the latest global status assessment issued by the UNEP [1], the energy requirements of the building sector experienced an approximate 4% expansion relative to 2020. The daily functioning of facilities is responsible for 30% [2] of worldwide energy use, with air-conditioning systems alone contributing almost 20% [3]. This demand trajectory is expected to continue upward, driven by elevating living standards and growing consumer expectations. Therefore, minimizing operational electricity usage while preserving occupant comfort remains a critical imperative for meeting aggressive sustainability targets. Roughly one-third of this expended energy stems from careless occupant practices, such as running air conditioners with open windows or selecting inappropriate thermostat parameters [4]. The automated identification of these occurrences offers a tremendous opportunity to unlock vast energy economies. Working under the assumption that energy conservation aligns with the Pareto principle (frequently referred to as the 80/20 rule), the bulk of these reductions can be achieved by isolating the irregular activities that represent just 20% of the total events. Such infrequent episodes, which exhibit abnormal usage signatures, are formally defined as anomalies.

This is an open access article distributed under the terms of the Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International License (CC BY-NC-SA 4.0), which permits non-commercial use, distribution, reproduction, and adaptation in any medium or format, provided that appropriate credit is given to the original author(s) and source, a link to the license is provided, and any changes made are indicated. Adapted material must be distributed under the same license.

© 2025 The Author(s). | Published by Institute of Smart Energy Systems Press

Tracking energy usage patterns represents the fundamental initial phase of anomaly identification. Contemporary infrastructure is equipped with smart metering technology to log power consumption. By compiling readings from every meter within a structure, a multivariate time series is formed, where each meter acts as an individual feature. The subsequent phase requires the formulation of a predictive model capable of characterizing the facility's normative usage profile, a baseline heavily dictated by environmental factors and human occupancy. Because these dynamics are typically non-linear, deep learning architectures are highly suitable for mapping them [5].

Based on the characteristics of the available data, these modeling frameworks fall into supervised, semi-supervised, or unsupervised categories. Compiling a comprehensive, fully labeled dataset that encompasses every conceivable real-world anomaly necessitates an impractical expenditure of expert human labor. Consequently, unsupervised algorithms present a highly viable path for practical deployment. The third stage of this pipeline entails generating a reconstructed usage profile via the model, which is then evaluated against the actual empirical measurements to quantify any discrepancies. The final step requires creating a specific evaluation metric to interpret this variance, ultimately allowing the system to flag aberrant data points as anomalies.

This research primarily addresses the second phase of the outlined workflow: engineering a computational model capable of mapping the complex interrelationships inherent in building energy usage data. Generally speaking, these modeling paradigms are divided into two distinct categories: predictor–detector strategies and pattern-based methods.

Predictor–detector-based approach

Within the predictor–detector framework, standard energy usage is essentially treated as a time series forecasting task. Numerous predictive algorithms, such as ARIMA, integrated ANN-ARIMA models [6], and deep Recurrent Neural Networks (RNNs) [7], are well documented in the academic literature. The projections generated by these models are subsequently routed into a detection module, which classifies a specific measurement as anomalous depending on its divergence from the actual recorded value. Researchers have frequently relied on straightforward, intuitive thresholds to evaluate this variance, such as exceeding twice the standard deviation [6] or moving beyond a specified percentage of the forecast [8]. In a related study, Fan *et al.* [9] deployed a Generalised Extreme Studentized Deviate(GESD) algorithm to isolate anomalies, utilizing an ensemble architecture that incorporated multiple linear regression (MLR), autoregressive integrated moving average (ARIMA), support vector regression (SVR), random forests (RF), multilayer perceptrons (MLP), boosting trees (BT), multivariate adaptive regression splines (MARS), and k-nearest neighbours (kNN). Comparatively, Xu and Chen [7] applied quantile regression to function as the anomaly detection mechanism alongside deep RNNs. Ultimately, these forecasting strategies project future consumption by extrapolating from historical records, exclusively factoring in periodic variables.

Pattern-based approach

Within pattern-based frameworks, specific characteristics of usage profiles are isolated to distinguish normal data points from anomalous ones. This classification encompasses unsupervised methods such as one-class classification, clustering, association rule mining, and dimensionality reduction [10]. For instance, Yu *et al*. [11] applied association rule mining to map temporal inter-feature correlations, establishing thresholds on these dependencies to detect abnormalities. While this framework offers strong interpretability, selecting the appropriate features demands significant domain knowledge. Alternatively, Arjunan *et al*. [12] proposed a dual-stage classifier that initially assigns scores to data instances and then employs a clustering mechanism to refine these scores using adjacent usage profiles. Although this dual-step strategy effectively isolates users exhibiting abnormal usage, it lacks the high interpretability inherent in rule-based models.

© 2025 The Author(s). | Published by Institute of Smart Energy Systems Press

Furthermore, clustering algorithms necessitate manual feature selection—a limitation that can be resolved through dimensionality reduction. Such reduction methods operate on the premise that most variables are superfluous, allowing a compact, robust feature set to be derived from the initial high-dimensional data. In this domain, autoencoders represent advanced tools for dimensionality reduction, particularly adept at modeling non-linear dependencies.

Fan *et al*. [13] evaluated the performance of multiple autoencoder architectures in the context of building energy anomaly detection, finding that a CNN-integrated autoencoder surpassed a standard ANN-based variant. Their research demonstrates that combining autoencoders with structures like CNNs effectively extracts temporal dynamics, making them highly viable for anomaly identification. However, such configurations still struggle to identify complex, multimodal interactions present across different variables. To address this, Guo *et al*. [14] introduced a GRU-based autoencoder focused on dimensionality reduction. This autoencoder is linked with a Gaussian Mixture Model (GMM) to map cross-feature interactions within the compressed latent space. While this strategy achieved superior performance compared to existing methods, it suffers from a lack of interpretability, as the complex associations between monitored variables remain opaque compared to rule-based alternatives. Additionally, the sequential data processing inherent to recurrent networks like GRUs can substantially escalate training duration when handling extensive multivariate time series datasets.

Consequently, unsupervised methodologies present a highly viable solution for building anomaly detection given the general absence of comprehensively labeled datasets. Among these, rule-based strategies provide excellent transparency but heavily rely on human specialists for feature engineering. Dual-stage methods—which combine historical scoring with localized neighborhood corrections—successfully map structural relationships but fall short in transparency. Similarly, autoencoders augmented with RNNs or CNNs, even when paired with clustering techniques, can isolate temporal patterns but remain largely uninterpretable. Therefore, a critical gap exists for an interpretable architecture capable of extracting both temporal dynamics and inter-variable correlations without demanding excessive training periods. **Table 1** contrasts the methodology proposed in this work against existing literature.

Aim, hypothesis and validation

This research endeavors to characterize normative energy usage profiles using an attention-augmented autoencoder, explicitly treating input variables as mutually dependent. **Figure 1** provides a visual summary of this framework. The core hypothesis posits that integrating inter-feature dependencies yields a superior reconstruction of the original data windows compared to assuming feature independence, particularly when the analyzed time window is shorter than the dataset's inherent cyclic periodicity. To clarify this premise, envision a facility containing multiple air-conditioning systems (e.g., one per occupant), each tracked via smart meters. Each individual's usage profile is influenced by both unique personal habits and shared environmental factors. While unique behaviors can be mapped using historical usage logs, shared influences necessitate a more nuanced analysis. Seasonal shifts are a prime example of shared factors. These variations can be captured by utilizing observation windows that span the duration of the cycle (often several months), which essentially localizes the effect to the specific user and justifies treating features independently. However, deploying extended time windows drastically inflates computational training requirements. Alternatively, by modeling features as interdependent, the system can successfully reconstruct usage patterns—even without fully capturing the extended temporal sequence—as long as correlations between features remain stable across the analyzed dataset windows. This technique bypasses the necessity for prolonged temporal inputs at the cost of mapping inter-feature linkages. Because these relationships are frequently time-variant and non-linear, this study leverages an attention mechanism [15]. Autoencoders are selected as the foundational architecture due to their proven capacity for dimensional compression, ensuring that only the most prominent characteristics are preserved during signal reconstruction, effectively screening out anomalous points.

© 2025 The Author(s). | Published by Institute of Smart Energy Systems Press

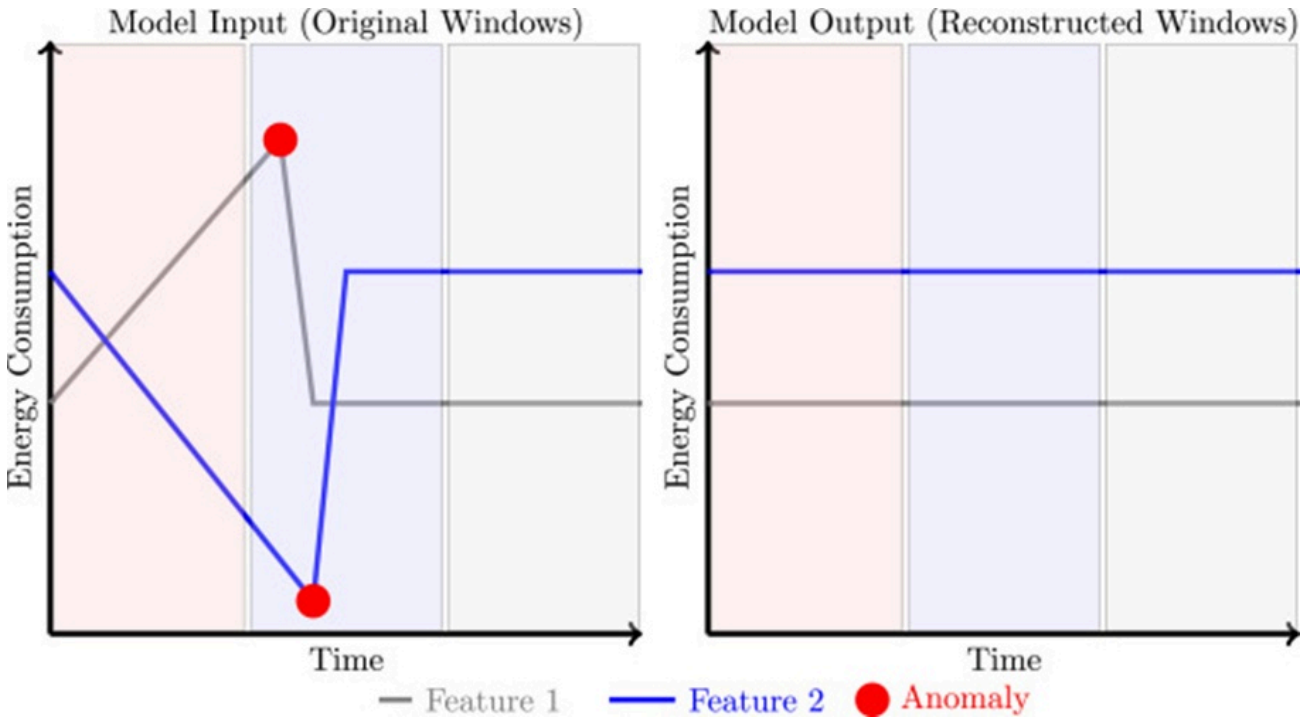

Figure 1. An illustration of the input and output of the model. Each box denotes data of size $M \times T \times N$, where $M = 3$ is the number of samples in each window and $N = 2$.

The suggested methodology undergoes validation across two scenarios using artificial datasets. Initially, the architecture is evaluated against a degenerate scenario—where the analysis window matches the data's maximum periodicity and features possess no mutual relationships—by benchmarking its performance against a standard ANN-based autoencoder. The objective here is to confirm whether the proposed framework matches the baseline in isolating temporal patterns. Following this, the model's reconstructive capacity is tested using time windows shorter than the dataset's overall maximum periodicity. Because this constraint risks generating artificial correlations, attention maps are deployed to ensure the underlying reconstruction logic matches the parameters required during data synthesis. Ultimately, the framework is applied to an empirical, real-world dataset, showcasing how attention maps can be utilized to decipher the internal mechanics of the model's signal reconstruction.

In summary, the main contributions of the study are:

1. Addition of the attention layer enables the autoencoder to weigh the input features to reconstruct windows adaptively

2. Demonstrating the application of attention to understand the reconstruction mechanics of the model

Table 1. Comparison between existing literature and proposed work. The columns indicate whether the technique incorporates those relations. T-T: Between two time windows, F-F: Between two features in a time window, I: Interpretability of the model, AFE: Automated Feature Engineering, Y: Available in the model and N: Not available in the model.

Methodology	T–T	F–F	I	AFE
Hybrid ANN integrated with ARIMA and 2 SD [6]	Y	N	N	N

© 2025 The Author(s). | Published by Institute of Smart Energy Systems Press

Deep RNN combined with Quantile Regression [7]	Y	N	N	N
Ensemble Framework employing GESD [9]	Y	N	N	N
Dual-Stage Classification Model [12]	Y	Y	N	N
CNN-Based Autoencoder Architecture [13]	Y	Y	N	Y
GRU Model coupled with GMM [14]	Y	Y	N	Y
Association Rule Mining Approach [11]	Y	Y	Y	N
Autoencoder enhanced with TiFe Attention Mechanism	Y	Y	Y	Y

Materials and Methods

Overview of problem

Let S represent a multivariate time series comprising N variables, where each variable reflects the measurements from an individual meter. This sequence S includes both typical and anomalous instances. By sampling S over time, observation windows of duration T are created. Each window is represented by a matrix X with dimensions $T \times N$. With M total windows available, the primary goal is to engineer a framework capable of isolating standard operational trends. During training, the framework distills essential characteristics from the dataset to optimize reconstruction accuracy. Subsequently, the architecture rebuilds the input window utilizing these identified variable interdependencies. A reconstruction loss metric is then computed, establishing a specific cutoff to flag data points exhibiting atypical behavior. For this study, we introduce an integrated approach that merges an attention module (designated as TiFe attention) with an encoder–decoder structure to facilitate sample reconstruction.

Description of datasets

This research utilizes three distinct datasets, detailed below:

1. Data-I is designed to validate the suggested framework's capacity to recognize temporal dynamics. It spans a full year and features a cyclic recurrence of weekly energy usage profiles **(Figure 2a).** Throughout the dataset, abnormal events are stochastically introduced across five weeks. An example of a standard anomalous sequence is depicted in **Figure 2b.**

© 2025 The Author(s). | Published by Institute of Smart Energy Systems Press

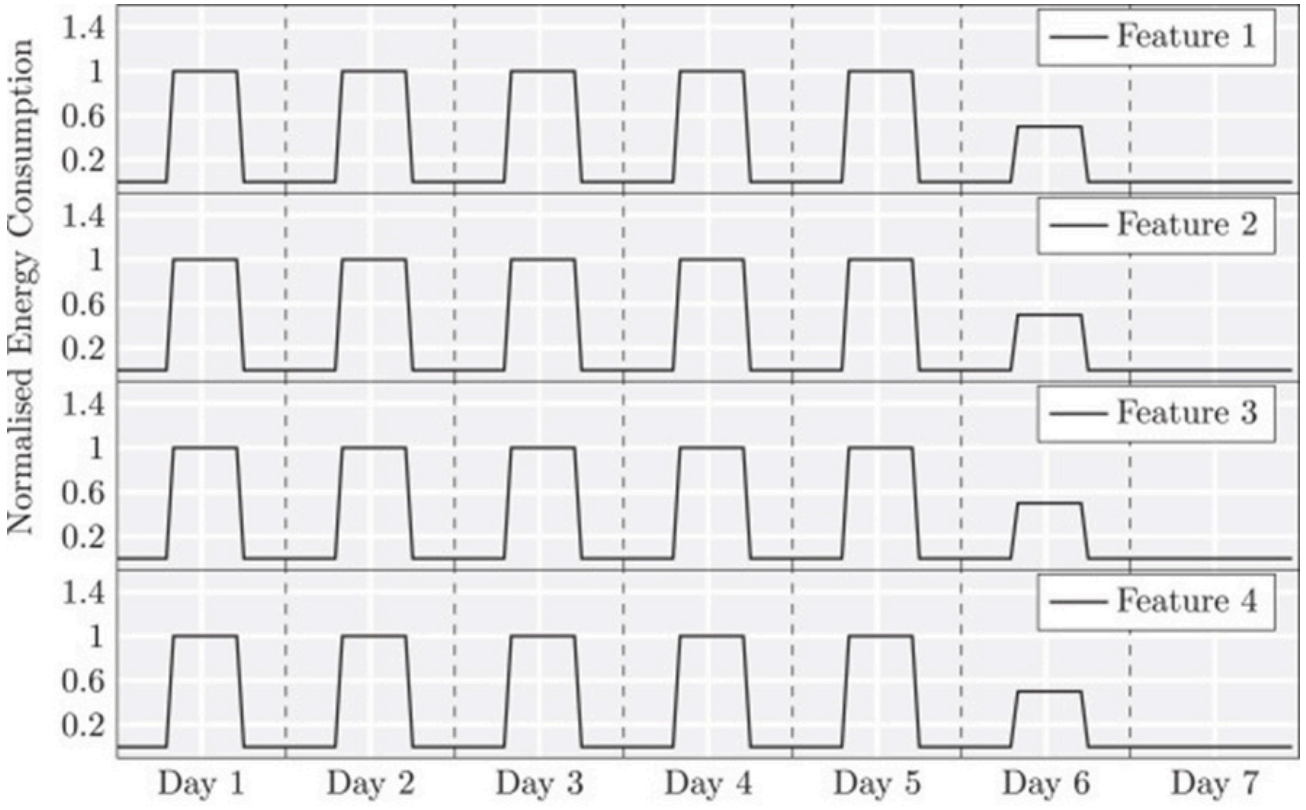

(a) Normal consumption

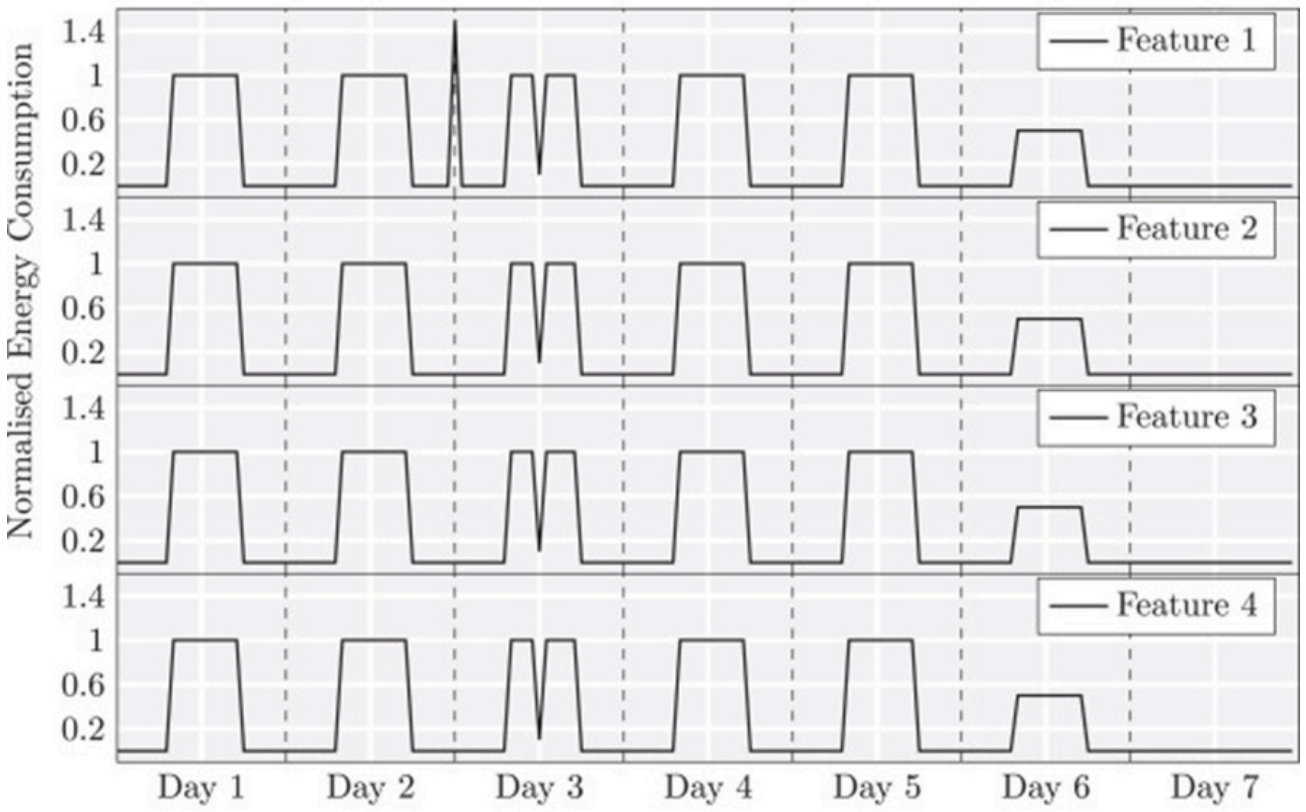

(b) Anomalous consumption

© 2025 The Author(s). | Published by Institute of Smart Energy Systems Press

Figure 2. Data-I: Energy consumption profiles for one week. Here $M = 1$, $T = 168(24 \times 7)$ and $N = 4$.

2. Data-II is intended to assess the architecture's effectiveness in mapping inter-variable relationships. This set comprises operational logs from two air-conditioning units managed by a central controller. This controller alternates the operation of the units every week, resulting in a dataset with a two-week cyclic period. Representative usage profiles are presented in **Figure 3.**

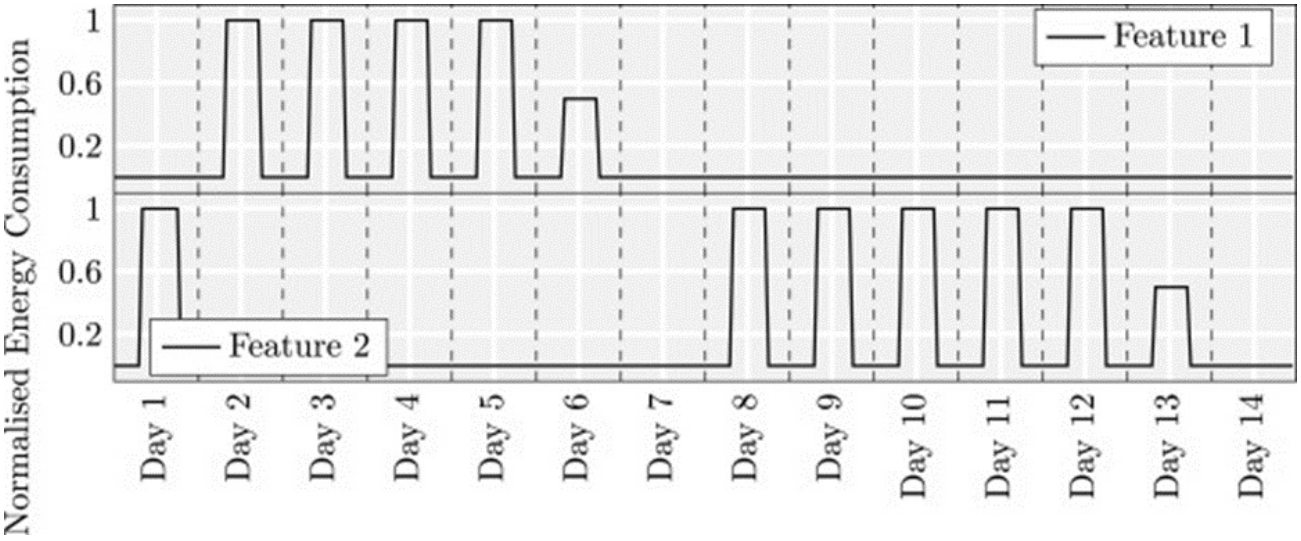

Figure 3. Data-II: Energy consumption profile for two consecutive weeks.

3. Data-III evaluates the framework's efficacy using an empirical, real-world dataset referenced in Pipattanasomporn *et al*. [16]. This collection contains two years of energy usage records from a university office building. Total power usage is categorized into lighting, plug loads, and HVAC. This study specifically isolates the HVAC component, encompassing 53 individual air-conditioning systems distributed across multiple zones on seven different floors. Preprocessing procedures involved min-max normalization and the imputation of absent data points using mean averages. This specific dataset was selected because the units exhibit interconnected spatial and temporal dynamics. Specifically, air-conditioners within a shared zone demonstrate strong correlations, whereas those in separate zones do not. The objective is to decode these interconnections utilizing attention maps.

Model architecture

Attention layer

Following their inception, LSTM [17] networks emerged as the standard approach for processing time series information. However, as recurrent structures, LSTMs analyze time windows sequentially, which can exponentially increase training time. This limitation is mitigated by employing a transformer architecture, as demonstrated in Vaswani *et al*. [15], which incorporates an attention mechanism allowing the network to evaluate data across all windows and detect underlying correlations simultaneously. Furthermore, transformer frameworks are inherently suited for parallel processing during training.

The fundamental principle of the attention mechanism involves formulating a structure of keys and values. When a query is introduced to this layer, it calculates a similarity metric (specifically, a dot product in this scenario; **Figure 4**) between the given query and the corresponding key. This metric is then normalized via a softmax function to generate attention

© 2025 The Author(s). | Published by Institute of Smart Energy Systems Press

scores, which subsequently modify the values before outputting the final vector. While this process has achieved benchmark performance in natural language processing, it is not directly transferable to building energy analytics. If an observation window $T \times N$ is input into a standard attention module, the associated queries, keys, and values are generated under the assumption that each column operates independently—a premise often invalid for building infrastructure data. To bypass this limitation, we introduce TiFe Attention, a dual-phase attention framework that evaluates both the rows and columns of the given window.

© 2025 The Author(s). | Published by Institute of Smart Energy Systems Press

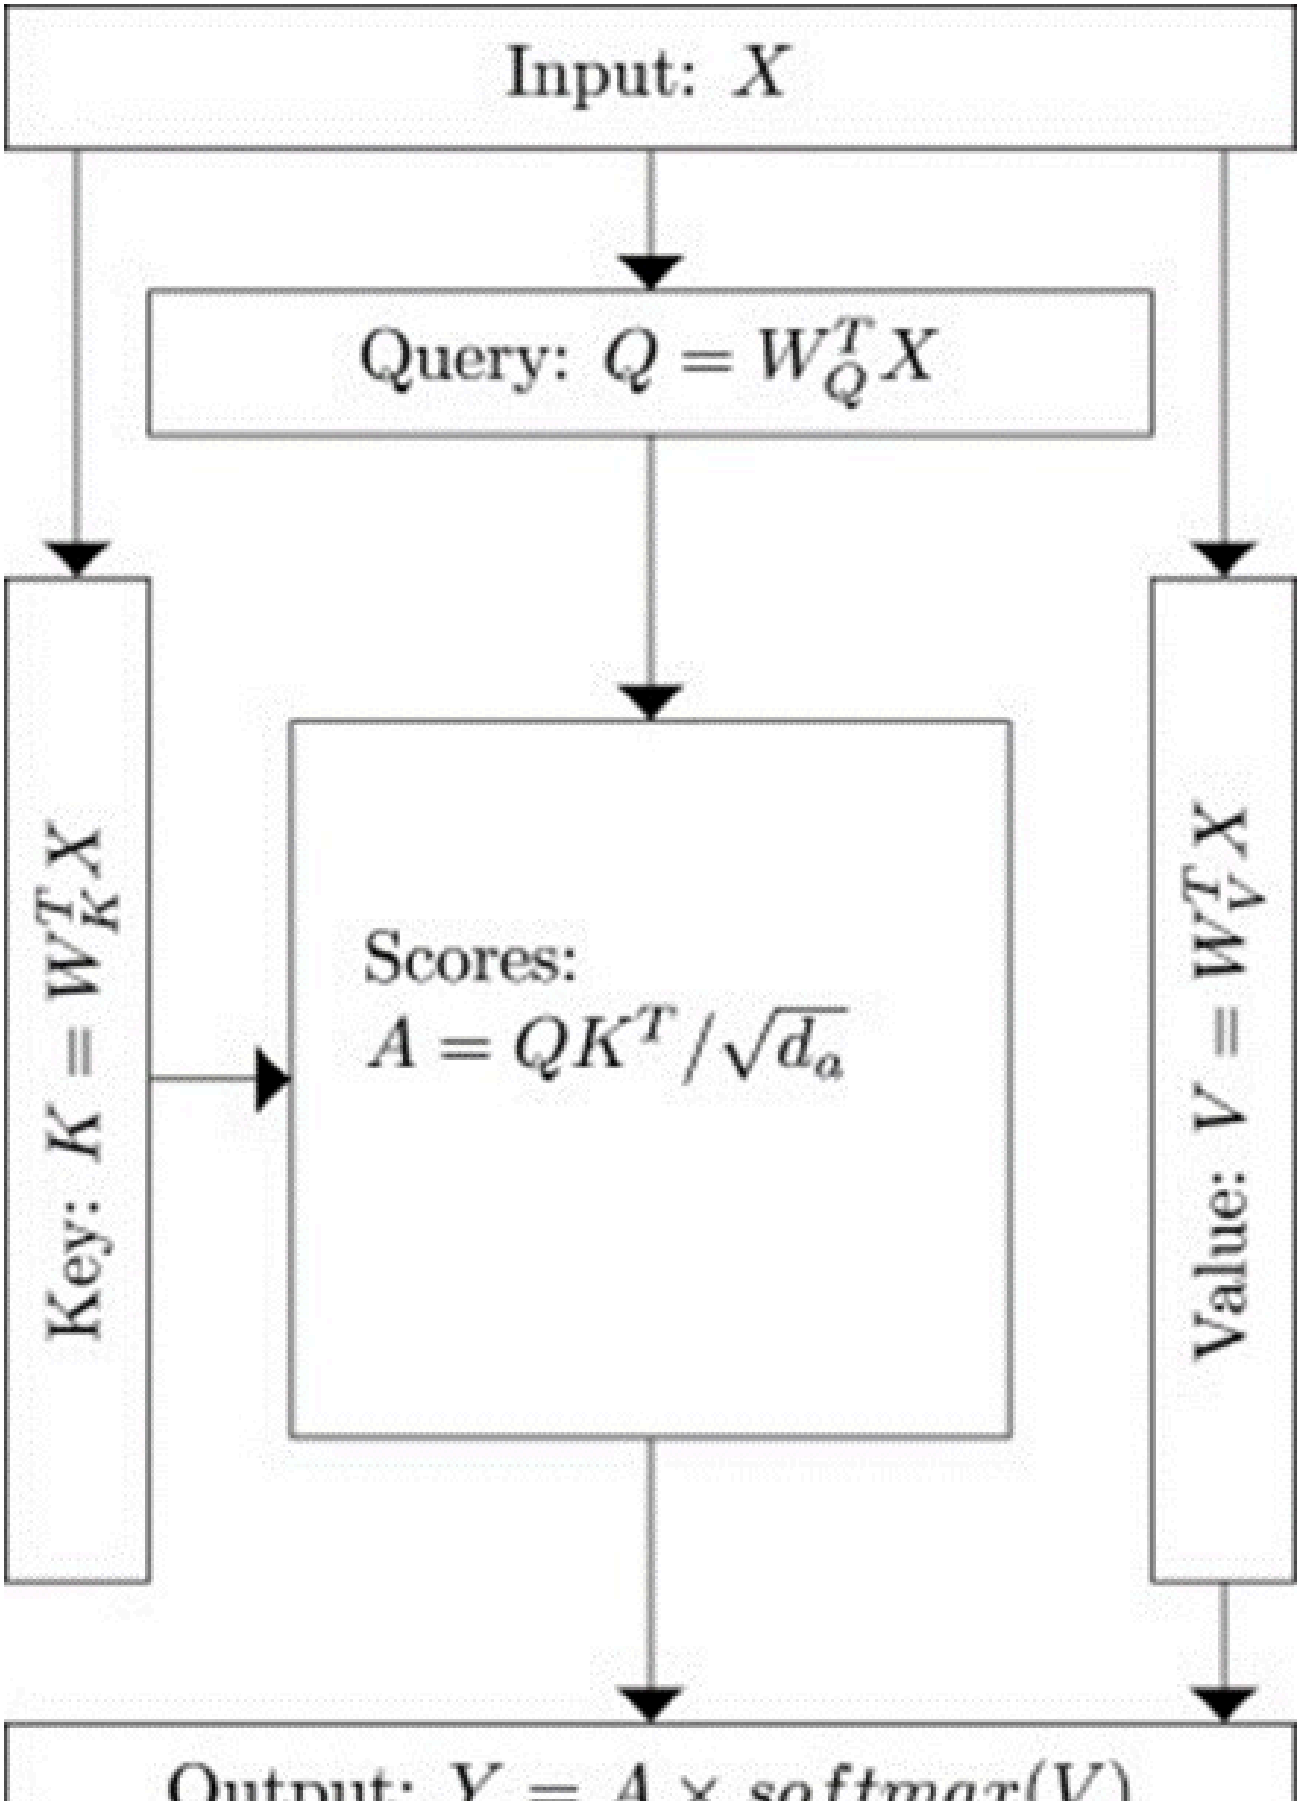

© 2025 The Author(s). | Published by Institute of Smart Energy Systems Press

ANNs.

As depicted in **Figure 5**, the TiFe attention module integrates an ANN extractor alongside two dot-product layers. The quantitative outputs derived from these attention modules are termed attention maps. Every attention map functions as a matrix, where the element located at the ith row and jth column represents the specific significance assigned to the jth variable during the reconstruction of the ith variable. The resulting feature and temporal attention maps are subsequently concatenated. Following this, the ANN extractor processes the merged attention outputs to generate the final output window. As illustrated in **Figure 6**, this transformation is executed through two sequential ANNs. Integrating this extractor ensures the architecture can identify cross-feature dependencies that may span different temporal durations. Nevertheless, decoding these complex relationships is impossible from a single attention map and necessitates examining maps across multiple temporal sequences. This study restricts its scope to evaluating attention maps derived from individual time windows.

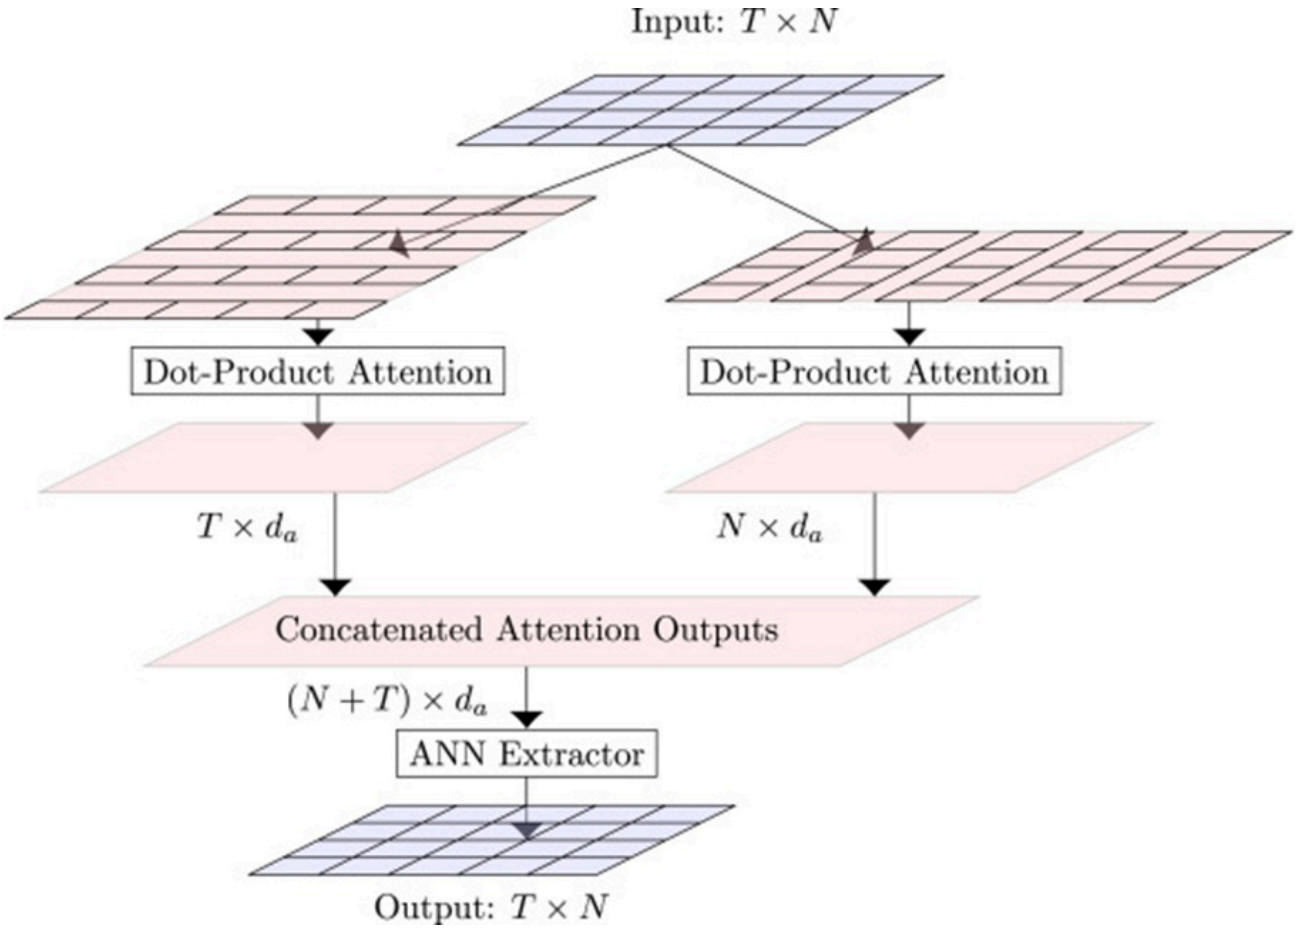

Figure 5. Illustration of the TiFe attention layer.

© 2025 The Author(s). | Published by Institute of Smart Energy Systems Press

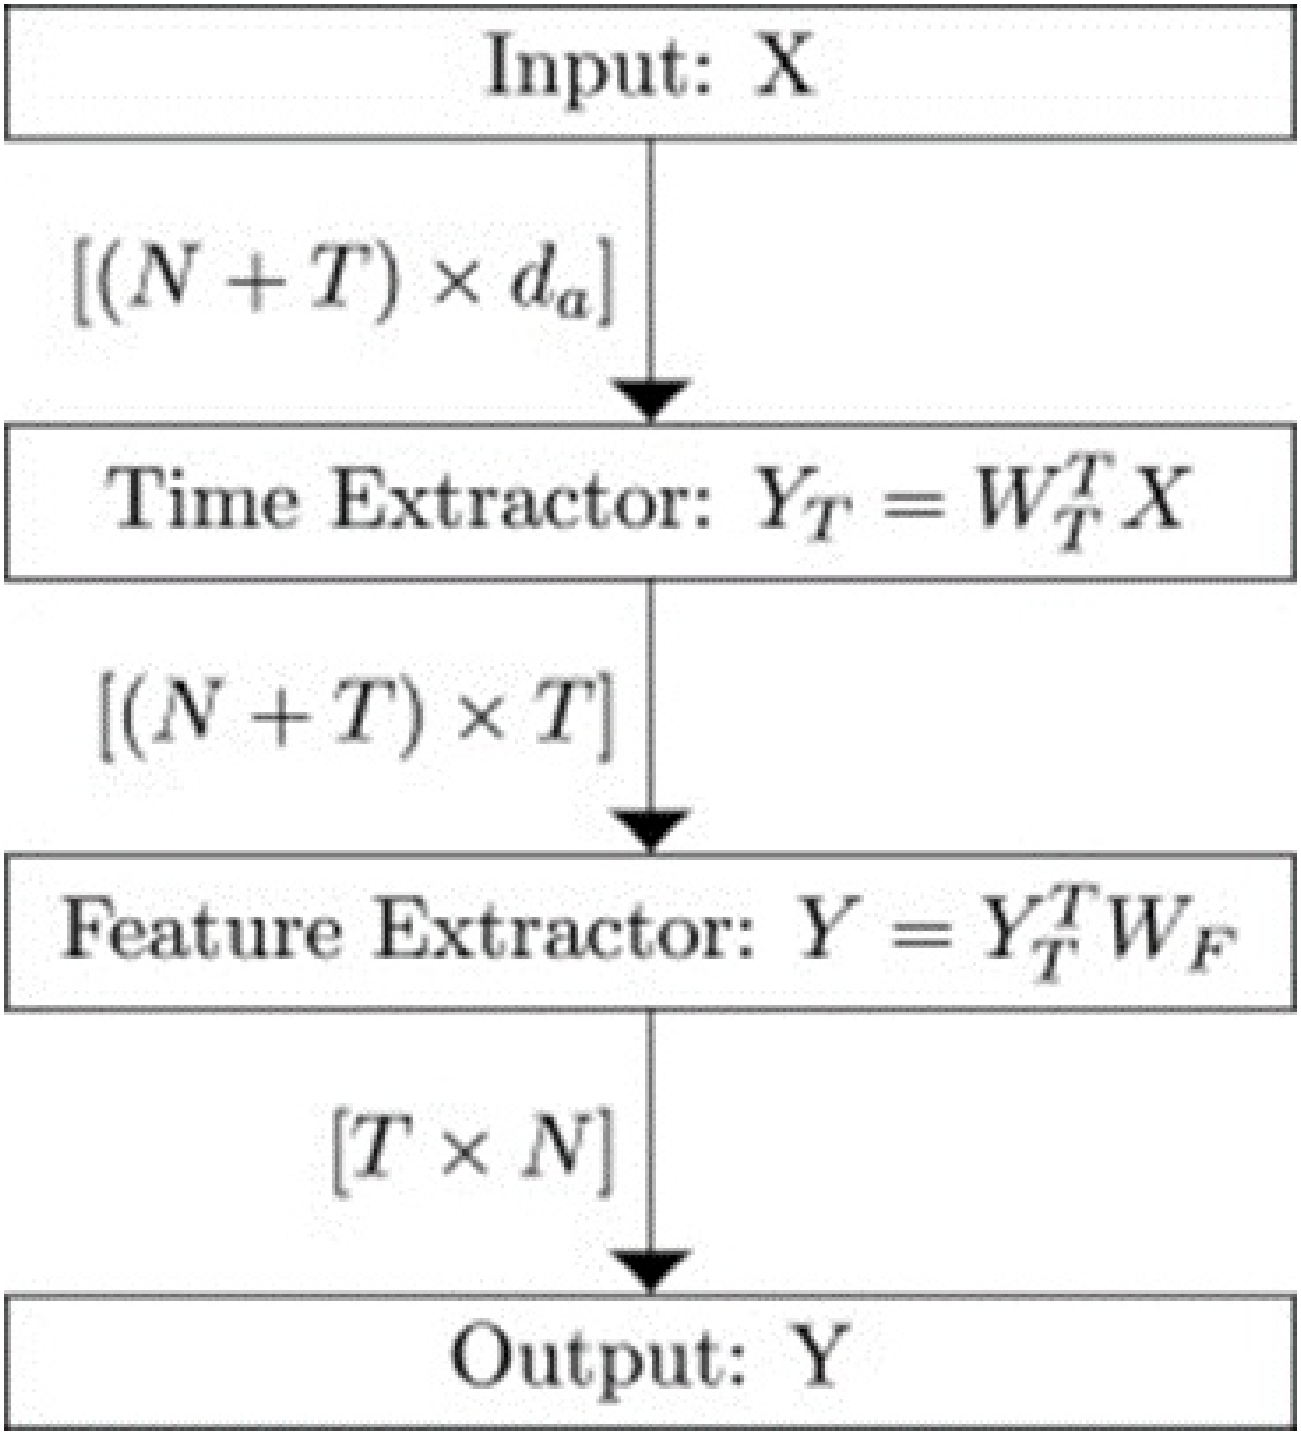

Figure 6. Illustration of the ANN extractor and the weights associated with the ANN in the transformation. is the number of units in each ANN.

Encoder-decoder

The architecture features an encoder $E : \Re^{T\ xN} \rightarrow \Re^{T\ xl}$ and a corresponding decoder $D : \Re^{T\ xl} \rightarrow \Re^{T\ xN}$ where l represents the framework's latent dimensionality; the encoder's function is to compress the provided inputs into a latent

© 2025 The Author(s). | Published by Institute of Smart Energy Systems Press

representation, whereas the decoder reverses this process, rebuilding the data from this latent state. In this study, ANNs function as the primary mathematical transformation operators. Given that the algorithm is trained exclusively on normal operational patterns, inputting an anomalous instance will inherently produce a significantly elevated reconstruction error. This resultant reconstruction loss acts as the primary indicator for anomaly detection. Typically, an arbitrary threshold is established, classifying any points exceeding this limit as abnormal. For the scope of this research, we introduce a qualitative method for interpreting this reconstruction discrepancy by leveraging the generated attention maps.

The complete structural layout of the proposed model is detailed in **Figure 7.** Within this design, attention modules are integrated into both the encoder and the decoder units. To illustrate the utility of this configuration, imagine analyzing one time window spanning Monday through Sunday, and a second window covering Wednesday through Tuesday. The encoder's attention mechanism effectively identifies a universal transformation that standardizes the input representation for both scenarios. Consequently, the decoder requires a complementary attention module to perform the reverse transformation accurately. In this capacity, the integrated attention layers function analogously to RNNs or CNNs, strategically weighting the provided time windows before isolating their most prominent characteristics.

© 2025 The Author(s). | Published by Institute of Smart Energy Systems Press

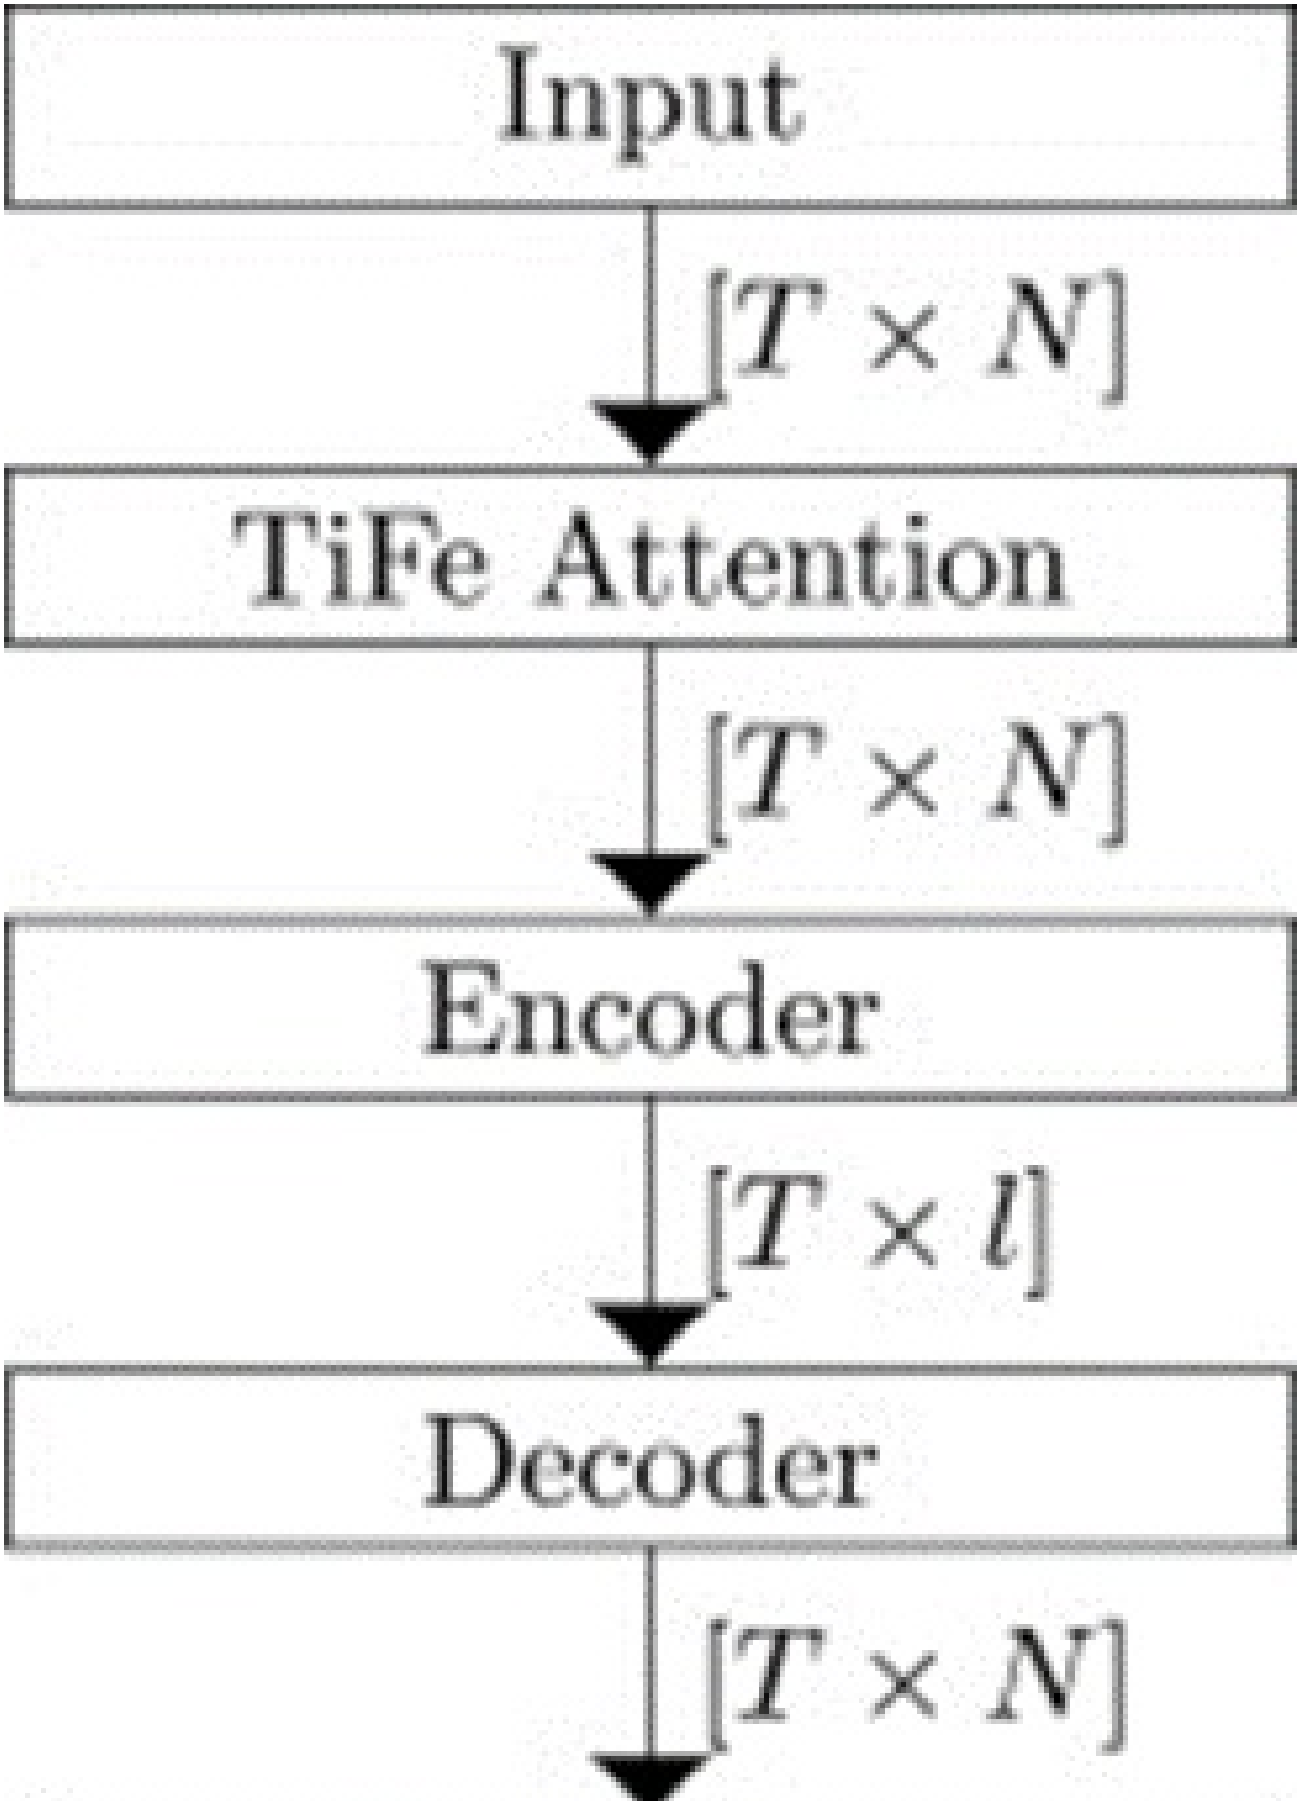

© 2025 The Author(s). | Published by Institute of Smart Energy Systems Press

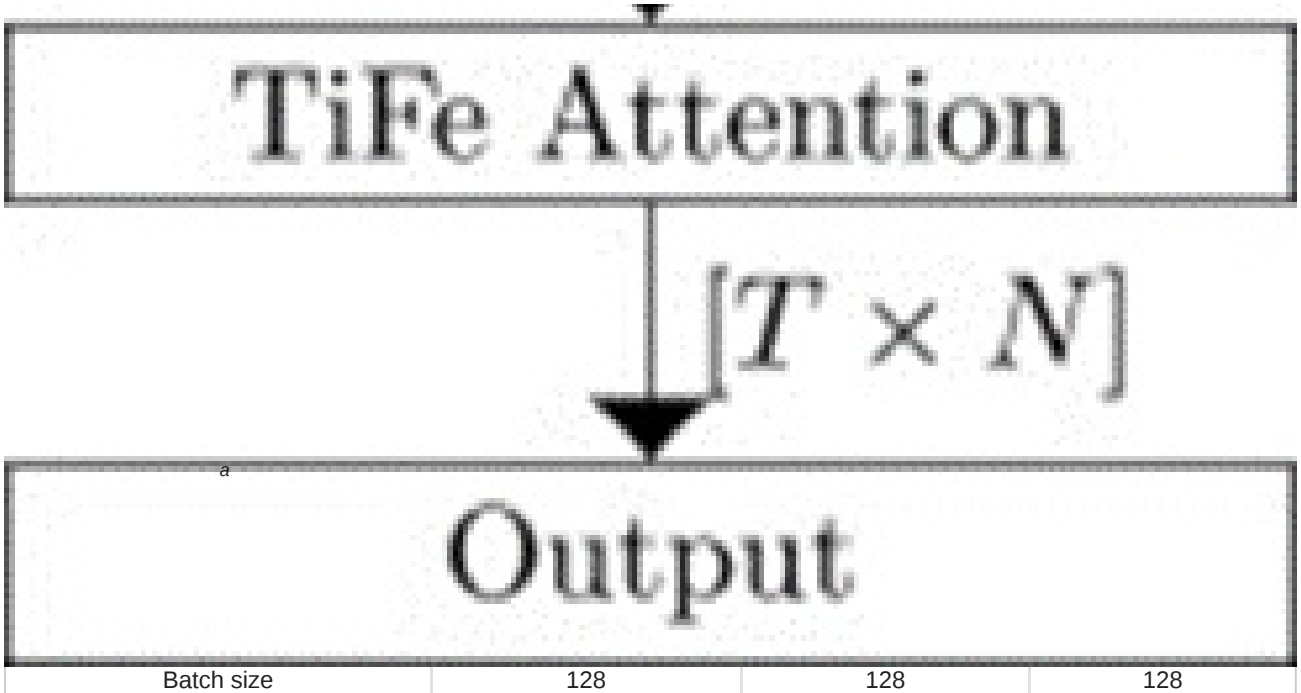

[a]

Batch size	128	128	128

Validation

Data-I evaluates the model's proficiency in rebuilding temporal sequences. Utilizing the proposed methodology, the framework is trained on anomalous data and evaluated against a baseline configuration featuring an ANN-based encoder-decoder system with an identical latent dimensionality. **Figure 8** illustrates the comparative reconstruction outcomes from both models. It is noteworthy that the peak at the conclusion of Day 2 and the trough during the middle of Day 3 are excluded from the TiFe attention reconstruction. Given that the attention layer represents the sole architectural distinction between the two, the capacity to capture the standard distribution accurately is directly linked to the TiFe attention framework. To comprehend how attention maps influence reconstruction, the temporal attention maps presented in **Figure 9** can be examined; a distinct band reflects the inherent daily cyclicity within the observation window, while the band's thickness denotes the duration of energy usage throughout the day. By contrasting the attention maps corresponding to both anomalous and normal instances, it becomes evident that the architecture dynamically adjusts the input weights before autoencoder processing, as depicted in **Figure 9b** for the late stages of Day 2 and the midpoint of Day 3. Consequently, the attention module adaptively calibrates the input windows to accurately simulate normative behavior.

© 2025 The Author(s). | Published by Institute of Smart Energy Systems Press

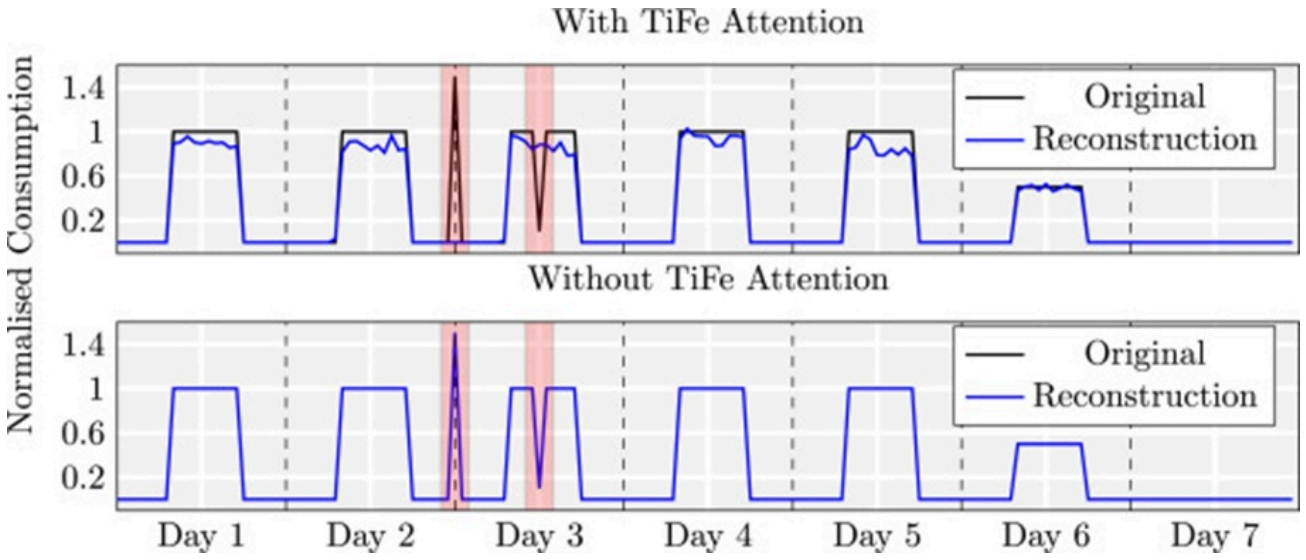

Figure 8. Data-I: Reconstruction comparison with and without TiFe. The highlighted regions indicate anomalies captured when using TiFe attention.

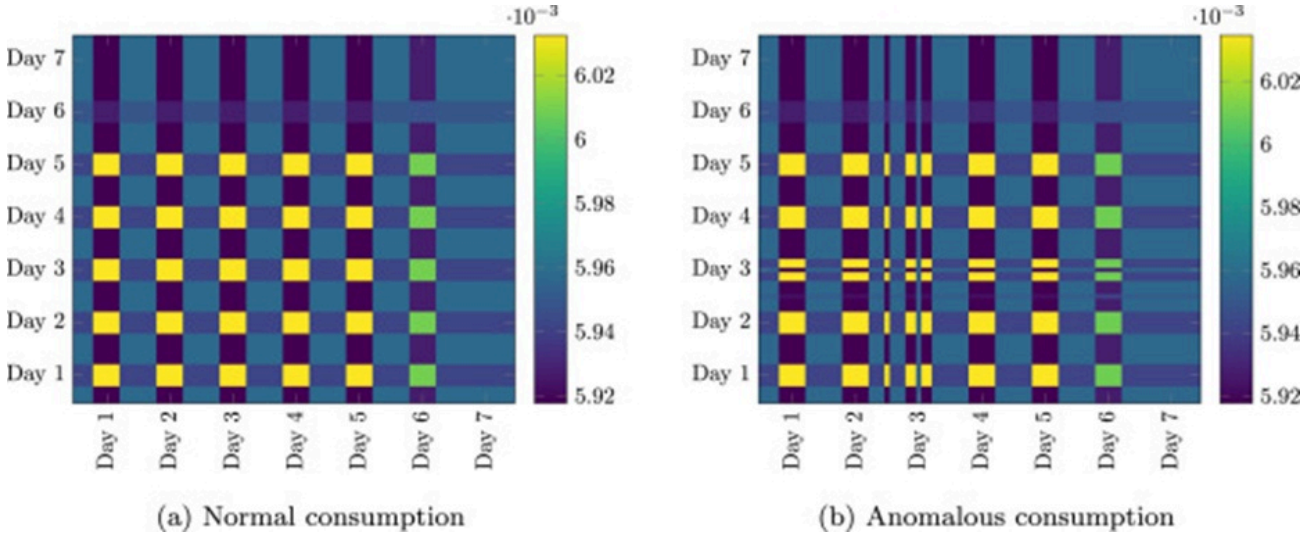

Figure 9. Data-I: Time attention maps indicating the inherent periodicity associated with data. Note the presence of additional lines in anomalous consumption.

Data-II validates the framework's capability to map inter-variable correlations. Because this dataset exhibits overlapping cyclic patterns (both daily and weekly) and the selected analysis window is shorter than the longest cyclical period, the algorithm is forced to discern the feature interdependencies that most accurately represent the underlying data. As a result, the training instances consist exclusively of either standard daily usage profiles or zero-consumption periods. Therefore, a basic autoencoder that treats features as isolated entities fails to accurately reconstruct the input window without anomalous artifacts accurately. Conversely, as demonstrated in **Figure 10**, the proposed architecture successfully executes the window reconstruction.

© 2025 The Author(s). | Published by Institute of Smart Energy Systems Press

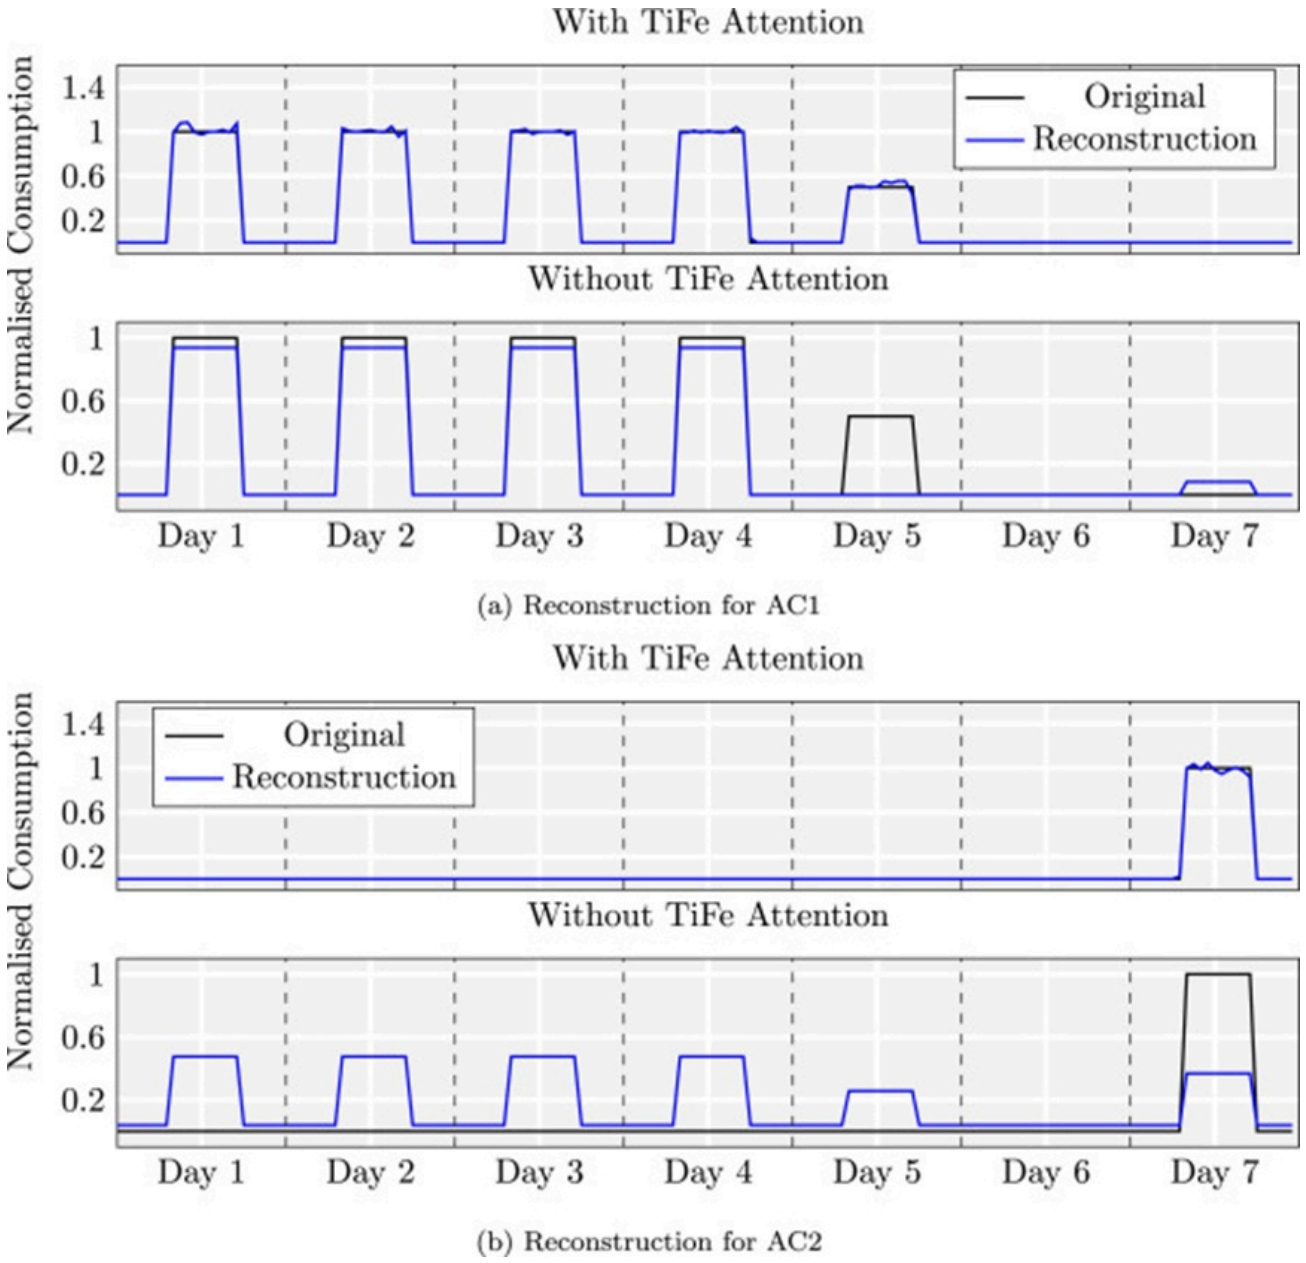

Figure 10. Data-II: Reconstruction comparison with and without TiFe for AC2.

Conversely, the TiFe-augmented autoencoder flawlessly reproduced the temporal windows across both variables. To further analyze the interplay between signal reconstruction and the generated attention maps, examining the feature attention map in **Figure 11b** is beneficial; the reconstruction of AC1 relies on weighted inputs derived from both air-conditioning units. This indicates that the architecture uses data from AC2 to synthesize AC1, and vice versa. Furthermore, a review of the temporal attention maps in **Figure 11a** reveals banding patterns consistent with those seen in Data-I. This confirms that the framework intelligently assigns weights to inputs originating from both variables to successfully reconstruct the entire observation window.

© 2025 The Author(s). | Published by Institute of Smart Energy Systems Press

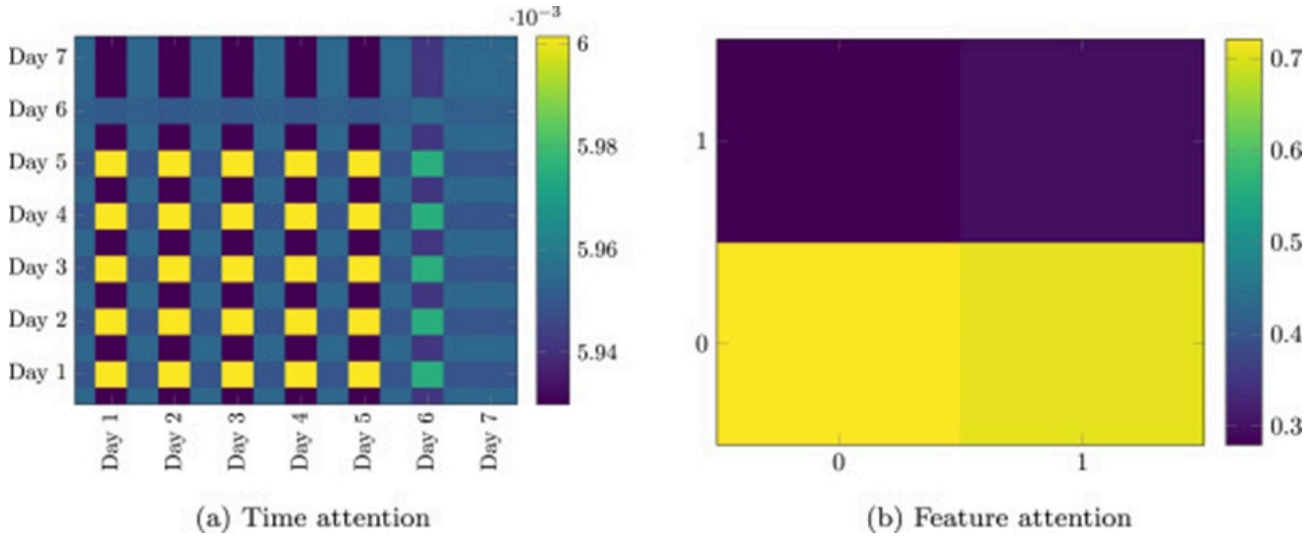

Figure 11. Data-II: Attention maps.

Results and Discussion

It is evident that the architecture successfully models both temporal dynamics and inter-feature interdependencies. Furthermore, attention maps prove to be a highly valuable mechanism for deciphering the internal logic the model employs to reconstruct a specific input sequence. Nevertheless, because the normative operational patterns in both Data-I and Data-II were already established a priori, the subsequent section illustrates the practical implementation of attention maps, detailing strategies to use them to tailor model training to the specific operational requirements of facility managers.

Analyzing the temporal attention map presented in **Figure 12**, the visible bands represent daily cyclical patterns. In contrast, the width of these bands corresponds to periods of active, non-zero power usage throughout the day. Examining the feature attention map in **Figure 13** reveals specific patches that highlight distinct variable clusters exhibiting analogous behavior. These groupings likely correspond to shared spatial zones or floors characterized by identical load profiles. To contextualize this, one can examine the reconstructed temporal windows for a specific identified cluster—namely features 10, 12, and 14—depicted in **Figures 14-16**, respectively. In each instance, Week 3 is fully reconstructed, despite presenting a usage profile that diverges noticeably from the prior weeks. Evaluating this reconstruction strictly through the lens of temporal distribution would yield an inaccurate assessment. However, the feature attention map **(Figure 13)** confirms a strong correlation among these variables, showing synchronous behavior throughout Week 3, which classifies this pattern as standard operation. This framework operates under the premise that anomalies represent only a marginal fraction of total observations, requiring the algorithm to reconstruct the vast majority of the dataset accurately. Because the framework is unsupervised, it simply derives an optimal set of inter-feature associations that adequately map the observed data during reconstruction. If these operational relationships are pre-established, the training phase can be explicitly bounded to adhere to these known dynamics. Consequently, facility operators can impose these predefined associations, thereby transitioning model training into a semi-supervised paradigm.

© 2025 The Author(s). | Published by Institute of Smart Energy Systems Press

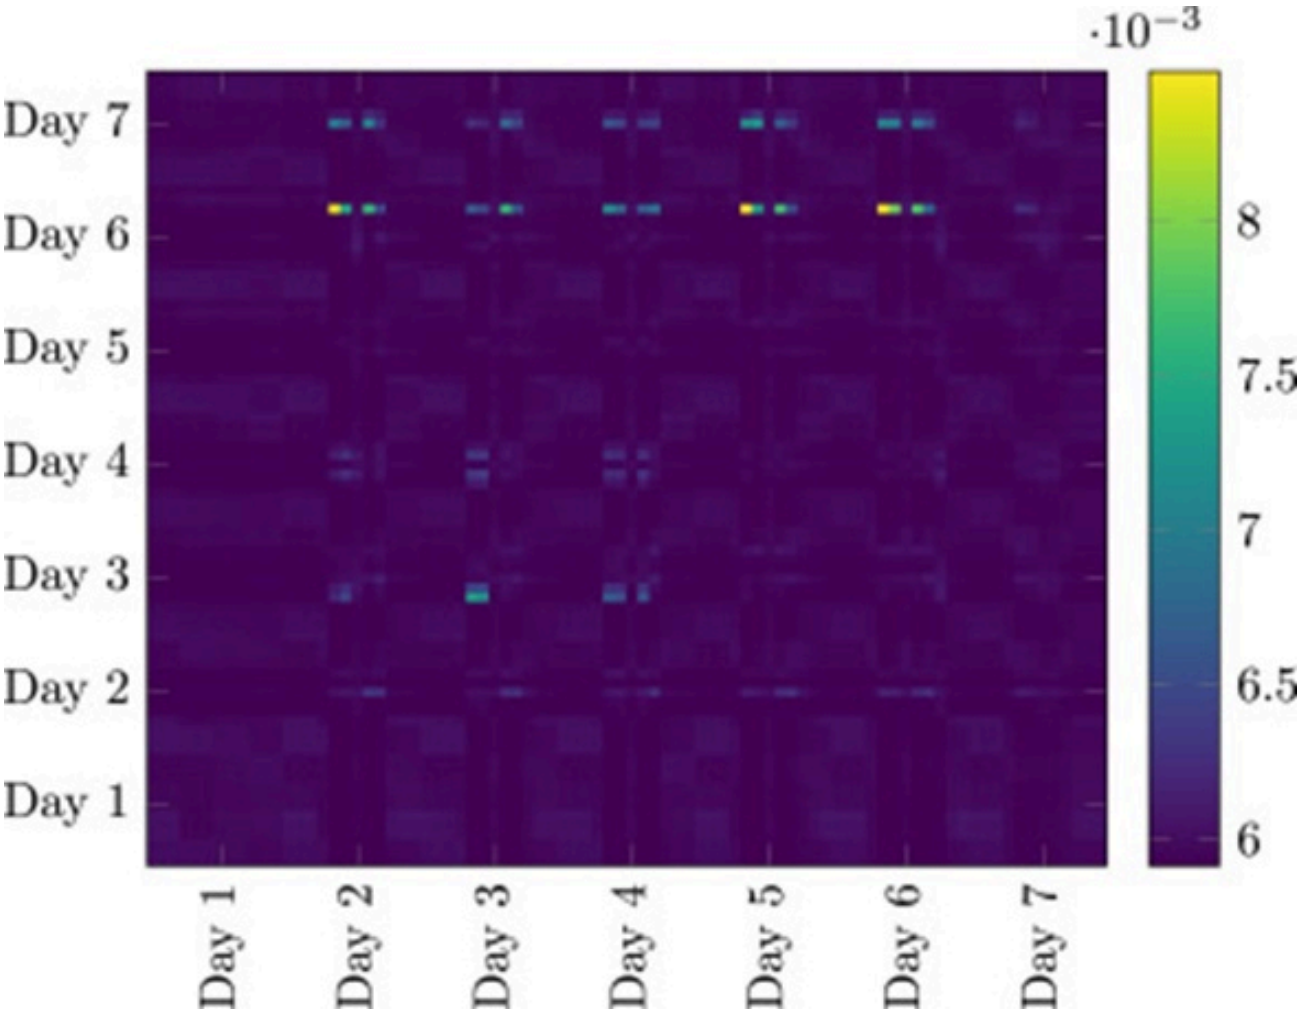

Figure 12. Data-III: Time attention map highlighting periodicity of the dataset.

© 2025 The Author(s). | Published by Institute of Smart Energy Systems Press

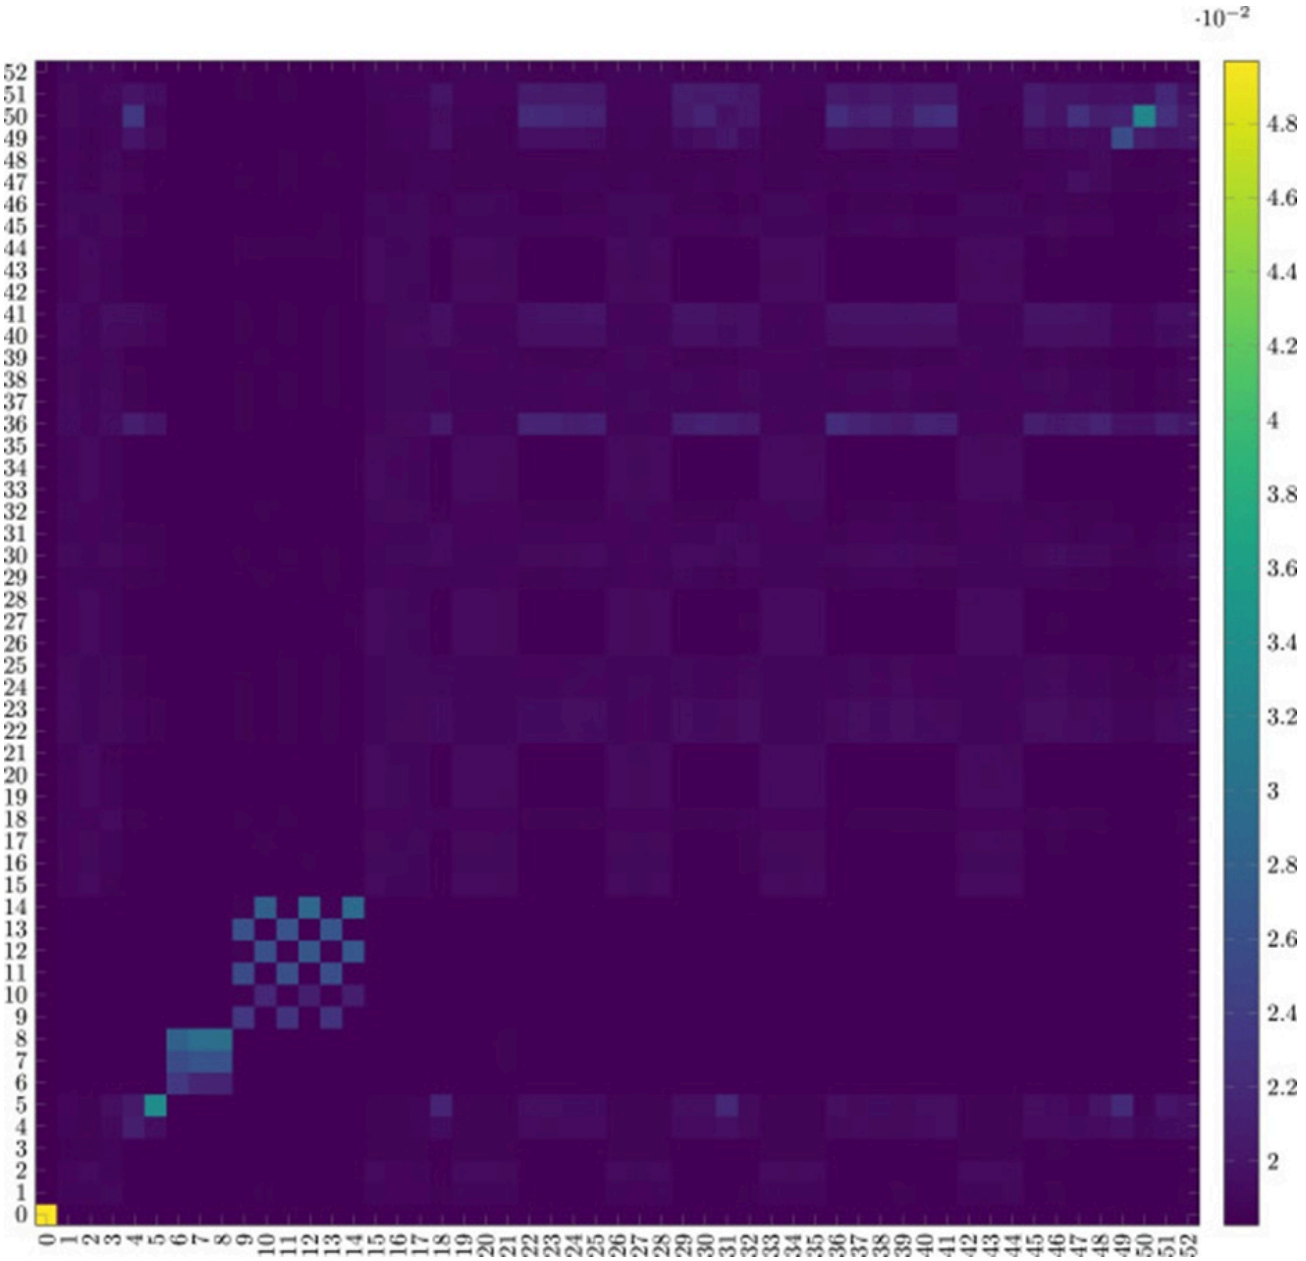

Figure 13. Data-III: Feature attention map highlighting the correlations among the features.

© 2025 The Author(s). | Published by Institute of Smart Energy Systems Press

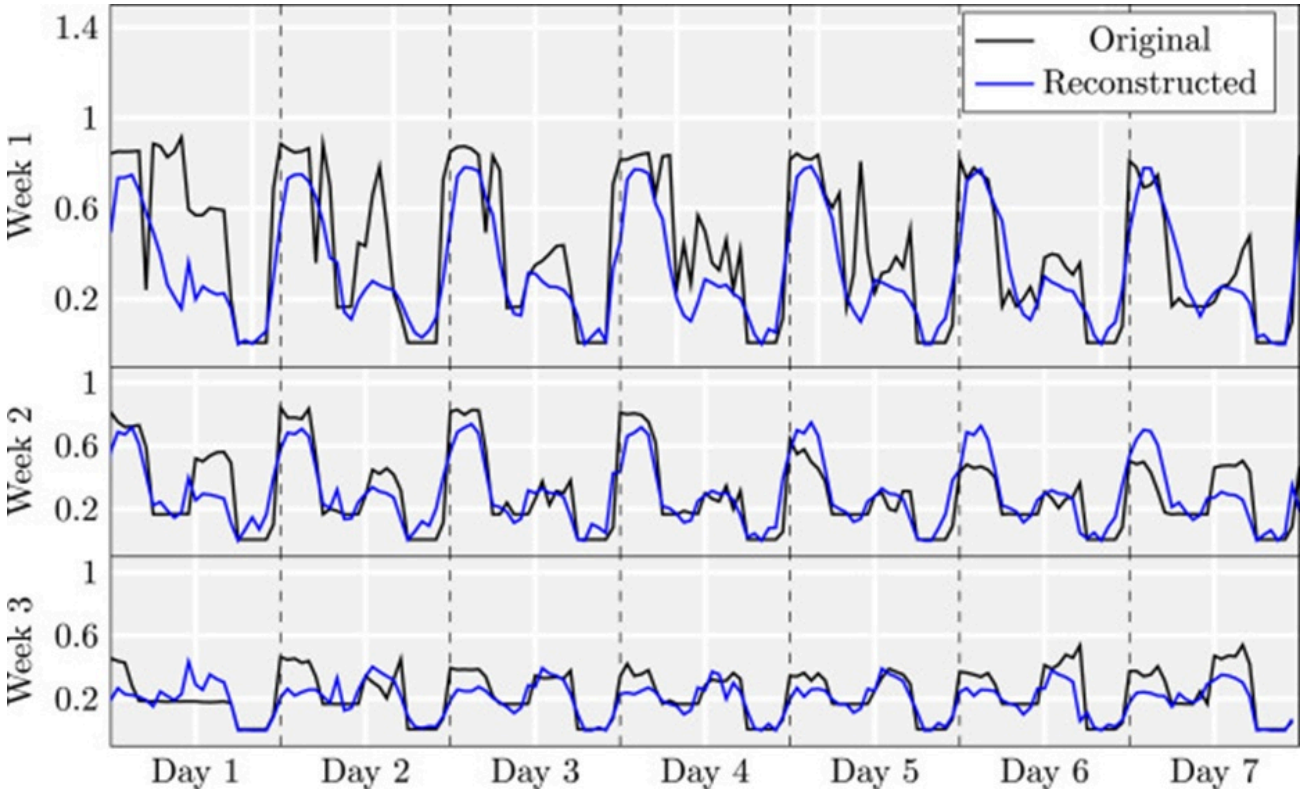

Figure 14. Data-III Reconstruction for three successive weeks for Feature-10. This corresponds to consumption from air-conditioning unit 6 on floor two and zone 2 of the building.

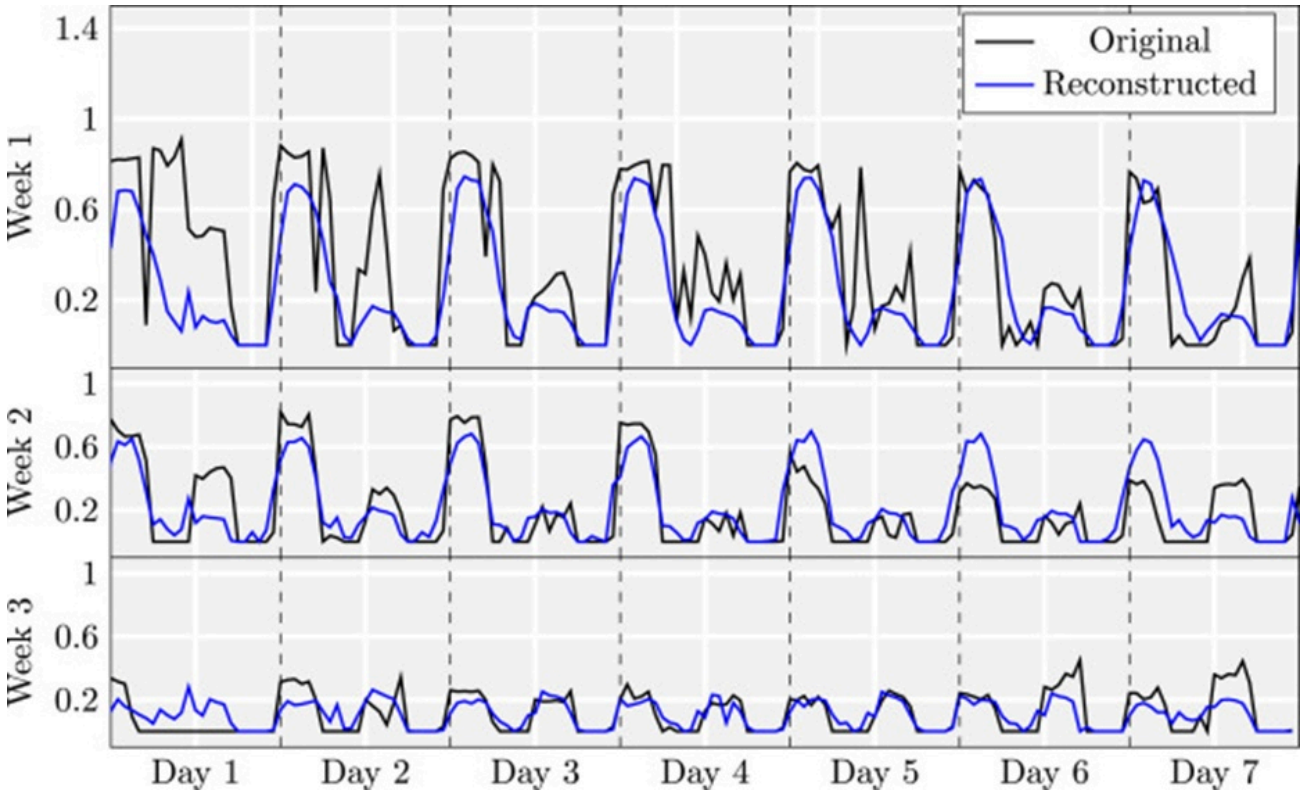

© 2025 The Author(s). | Published by Institute of Smart Energy Systems Press

Figure 15. Data-III Reconstruction for three successive weeks for Feature-12. This corresponds to consumption from air-conditioning unit 8 on floor two and zone 2 of the building.

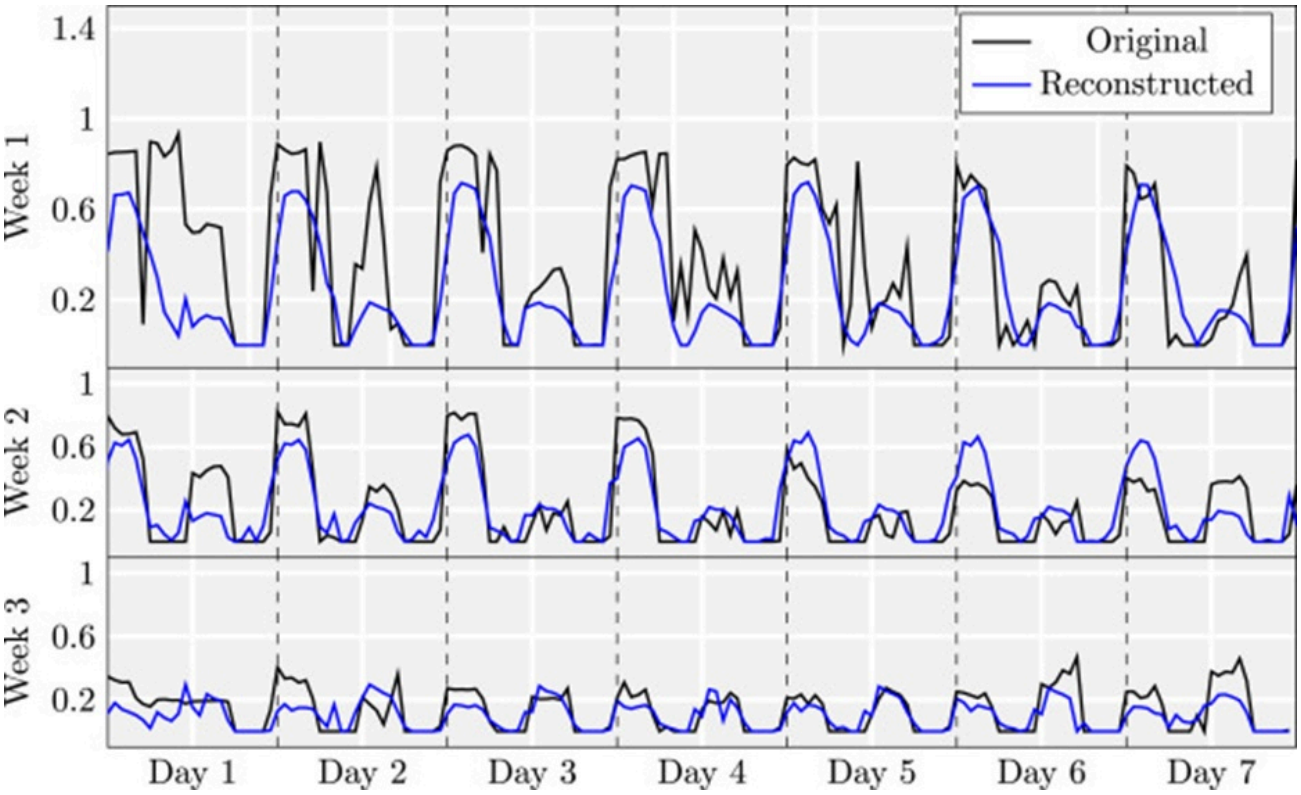

Figure 16. Data-III Reconstruction for three successive weeks for Feature-14. This corresponds to consumption from air-conditioning unit 10 on floor two and zone 2 of the building.

Conclusion

Implementing strategic energy management within facility operations offers substantial benefits for broader energy conservation initiatives. The automated identification of irregular power usage can facilitate the formulation of protocols designed to mitigate careless occupant habits. The fundamental principle driving this anomaly detection framework centers on mapping normative energy usage baselines and triggering alerts when deviations occur. While numerous machine learning algorithms exist to characterize these standard usage profiles, this study introduces an attention-based architecture specifically designed to map typical consumption dynamics. The framework's evaluation using artificially generated data validates its capacity to extract these essential patterns. Consequently, the derived attention maps provide a mechanism to optimize network parameters without succumbing to overfitting, even when the training corpus contains a mixture of standard and aberrant instances. However, the introduced framework is constrained by its core premise—that anomalies represent an exceedingly small fraction of the training data—and by the risk that a severely restricted latent space might compel the network to map superficial, non-physical correlations. To overcome these limitations, future research will explore implementing a semi-supervised modeling strategy.

Abbreviations	
ANN	Artificial Neural Network
ARIMA	Autoregressive Integrated Moving Average
BT	Boosting Tree

© 2025 The Author(s). | Published by Institute of Smart Energy Systems Press

CNN	Convolutional Neural Network
GESD	Generalised Extreme Studentized Deviate
GMM	Gaussian Mixture Model
GRU	Gated Recurrent Unit
HVAC	Heating, Ventilation and Air Conditioning
kNN	k Nearest Neighbors
LSTM	Long-Short Term Memory
MARS	Multivariate Adaptive Regression Splines
MLP	Multilayer Perceptron
RF	Random Forest
RNN	Recurrent Neural Network
SD	Standard Deviation
SVR	Support Vector Regression

Acknowledgements

None

Conflict of interest

None

Financial support

This research was supported by the Prime Minister's Research Fellows (PMRF) [project ID: SB22230925MEPMRF008998] under the Ministry of Education, Government of India.

Ethics statement

None

Received: 28 Jun 2023 Revised: 02 Sep 2023 Accepted: 18 Sep 2023
Published online: 15 June 2025

Rights and permissions

© 2025 The Author(s). | Published by Institute of Smart Energy Systems Press

Open Access The author(s) retain copyright. This article is licensed under the Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International License. It may be shared and adapted for non-commercial purposes with appropriate attribution, an indication of changes, and distribution of adaptations under the same license. Third-party material may be subject to separate terms identified in its credit line. View the license at https://creativecommons.org/licenses/by-nc-sa/4.0/.

References

1. UN Environment Programme. 2022 Global status report for buildings and construction: Towards a zero–emission, efficient and resilient buildings and construction sector. 2022.

2. IEA. Buildings. Paris: IEA; 2022 [cited 2026 Jul 6]. Available from: https://www.iea.org/reports/buildings

3. IEA. The future of cooling. Paris: IEA; 2018 [cited 2026 Jul 6]. Available from: https://www.iea.org/reports/the-future-of-cooling

4. Rashid H, Singh P. Monitor: An abnormality detection approach in buildings energy consumption. In: 2018 IEEE 4th International Conference on Collaboration and Internet Computing (CIC). IEEE; 2018. p. 16-25. https://doi.org/10.1109/CIC.2018.00-44.

5. LeCun Y, Bengio Y, Hinton G. Deep learning. Nature. 2015;521(7553):436-44. https://doi.org/10.1038/nature14539.

6. Sheng CJ, Telaga AS. Real-time detection of anomalous power consumption. Renew Sustain Energy Rev. 2014;33:400-11. https://doi.org/10.1016/j.rser.2014.01.088.

7. Xu C, Chen H. A hybrid data mining approach for anomaly detection and evaluation in residential buildings energy data. Energy Build. 2020;215:109864. https://doi.org/10.1016/j.enbuild.2020.109864.

8. Mavromatidis G, Acha S, Shah N. Diagnostic tools of energy performance for supermarkets using Artificial Neural Network algorithms. Energy Build. 2013;62:304-14. https://doi.org/10.1016/j.enbuild.2013.03.020.

9. Fan C, Xiao F, Wang S. Development of prediction models for next-day building energy consumption and peak power demand using data mining techniques. Appl Energy. 2014;127:1-10. https://doi.org/10.1016/j.apenergy.2014.04.016.

10. Himeur Y, Ghanem K, Alsalemi A, Bensaali F, Amira A. Artificial intelligence based anomaly detection of energy consumption in buildings: A review, current trends and new perspectives. Appl Energy. 2021;287:116601. https://doi.org/10.1016/j.apenergy.2021.116601.

11. Yu ZJ, Haghighat F, Fung BC, Zhou L. A novel methodology for knowledge discovery through mining associations between building operational data. Energy Build. 2012;47:430-40. https://doi.org/10.1016/j.enbuild.2011.12.018.

12. Arjunan P, Khadilkar HD, Ganu T, Charbiwala ZM, Singh A, Singh P, et al. Multi-user energy consumption monitoring and anomaly detection with partial context information. In: Proceedings of the 2nd ACM International Conference on Embedded Systems for Energy-Efficient Built Environments. 2015. p. 35-44. https://doi.org/10.1145/2821650.2821662.

13. Fan C, Xiao F, Zhao Y, Wang J. Analytical investigation of autoencoder-based methods for unsupervised anomaly detection in building energy data. Appl Energy. 2018;211:1123-35. https://doi.org/10.1016/j.apenergy.2017.12.005.

14. Guo Y, Liao W, Wang Q, Yu L, Ji T, Li P. Multidimensional time series anomaly detection: A gru-based gaussian mixture variational autoencoder approach. In: Asian Conference on Machine Learning. PMLR; 2018. p. 97-112 [cited 2026 Jul 6]. Available from: https://proceedings.mlr.press/v95/guo18a.html

15. Vaswani A, Shazeer N, Parmar N, Uszkoreit J, Jones L, Gomez AN, et al. Attention is all you need. Adv Neural Inf Process Syst. 2017;30. https://doi.org/10.48550/ARXIV.1706.03762.

16. Pipattanasomporn M, Chitalia G, Songsiri J, Aswakul C, Pora W, Suwankawin S, et al. CU-BEMS, smart building electricity consumption and indoor environmental sensor datasets. Sci Data. 2020;7(1):241. https://doi.org/10.1038/s41597-020-00582-3.

17. Hochreiter S, Schmidhuber J. Long short-term memory. Neural Comput. 1997;9(8):1735-80. https://doi.org/10.1162/neco.1997.9.8.1735.

© 2025 The Author(s). | Published by Institute of Smart Energy Systems Press