

# Development of a Decision Support System for Building Energy Management

Teodor-Lucian Ene  
Automation and Industrial Informatics  
University Politehnica of Bucharest  
Bucharest, Romania  
teodor\_lucian.ene@stud.acs.upb.ro

Grigore Stamatescu  
Automation and Industrial Informatics  
University Politehnica of Bucharest  
Bucharest, Romania  
grigore.stamatescu@upb.ro

**Abstract**—This paper presents the development stages and implementation results of a decision support system (DSS) for the energy management of office buildings. The research uses the Matrix Profile (MP) algorithm to detect anomalies in the Building Data Genome Project 2 dataset. The results demonstrate the system’s ability to identify HVAC faults and quantify energy waste without manual data labeling. The implementation was able to identify a saving opportunity of €29,800/year by optimizing equipment start-up (“Soft Start”) and a 15% reduction in residual consumption (“Phantom Load”). The system also calculated a carbon footprint of 2,044 tonnes of CO<sub>2</sub> using a conversion factor of 0.25 kg/kWh, providing a solid basis for Environmental, Social and Governance (ESG) reporting through an open software.

**Index Terms**—data mining, energy efficiency, matrix profile, decision support system, building automation.

## I. INTRODUCTION

The commercial building sector is a major consumer of energy globally, with a significant portion of this consumption attributable to operational inefficiencies, incorrect building management system (BMS) settings, or equipment failures that go unnoticed. In current practice, facility managers often rely on reactive analytics or static alert thresholds, which generate many false alarms and do not take into account operational or climatic contexts.

The objective of this research is to develop a functional prototype capable of automating the data analysis and decision support process in smart buildings. The goal was to move from a simple data visualization to a system capable of diagnosing the condition of the building and proposing concrete, financially quantified actions.

This article provides three main contributions to the field of building energy analysis: first, it introduces an unsupervised detection methodology by applying the Matrix Profile algorithm to eliminate the need for manually trained and labeled datasets, which are often not available in existing buildings; second, it provides thermodynamic contextualization by integrating the building energy signature with meteorological data to reduce the false alarm rate caused by extreme weather conditions; and third, it operationalizes the impact by transforming technical anomalies into financial and sustainability indicators that can be directly used in the managerial decision-making process.

The rest of this paper is organized as follows: Section II reviews the current state of building energy management and the theoretical foundation of the Matrix Profile algorithm. Section III describes the methodology and algorithms used for data selection and anomaly detection. Section IV presents the software implementation and its modular architecture. Section V discusses the experimental results, including the comparative analysis an impact of building selection and the quantified impact of the detected inefficiencies. Finally, Section VI concludes the paper and presents future research directions.

## II. STATE OF THE ART

Energy performance management in the built environment has undergone major transformations in the last decade, moving from sporadic manual checks to continuous digital monitoring. The widespread implementation of Advanced Metering Infrastructure (AMI) and IoT sensors has generated an exponential increase in the amount of data available to building managers. However, simply collecting these time series does not guarantee consumption optimization. The industry faces the difficulty of extracting useful information from raw data, a phenomenon often described as “data rich, information poor”. Currently, the specialized literature and industrial research directions focus on the development of automated software solutions capable of offloading the analysis burden from the human operator and moving from reactive to predictive maintenance [1].

### A. Evolution of Fault Detection and Diagnostics (FDD) Systems

The field known as FDD (Fault Detection and Diagnosis) has become a central component in intelligent building management strategies. The objective of these systems is to automatically identify suboptimal equipment operation before it generates additional costs or thermal discomfort. Katipamula and Brambley [2] classify FDD methods into three distinct methodological categories, each with specific advantages and limitations:

- 1) *Quantitative model-based methods*: This approach involves constructing a detailed mathematical model [3] of the building, based on thermodynamic and heat

transfer principles. Although theoretically offering the highest accuracy, their practical applicability is limited. Implementation requires detailed knowledge of building parameters, such as wall thermal resistance and equipment coefficients, information that is often unavailable or inaccurate for existing buildings. Calibrating these models is a laborious and expensive process, difficult to scale across a large real estate portfolio

- 2) *Knowledge-based qualitative methods*: This is the dominant approach in commercial solutions today and is based on sets of logical rules derived from engineering expertise (e.g., if-then logic). Expert systems verify compliance with predefined operating conditions. However, recent studies, including those by Granderson et al. [4], indicate that systems based strictly on static rules tend to generate a large number of false alarms when operating conditions change. This instability often leads to operators ignoring alerts (“alarm fatigue”), reducing the effectiveness of the system.
- 3) *Data-driven methods*: This category, which includes artificial intelligence and advanced statistical techniques, has gained significant ground. Unlike the first two approaches, these methods do not require physical knowledge of the building or explicit rules, but learn behavioral patterns directly from historical time series. This flexibility makes them suitable for rapid deployment in buildings with heterogeneous infrastructure, where data is abundant but technical information about equipment is lacking.

### *B. Existing Software Architectures and Limitations of Deterministic Approaches*

The embodiment of the building automation layer is mainly represented by the Building Management Systems (BMS) as an operational system, ensuring the operation of electromechanical equipment through continuous control and feedback loops. Although modern BMSs have the ability to store historical operating parameters, their architecture is optimized for immediate execution, rather than retrospective analysis. Consequently, alerts generated by a BMS are usually binary and are triggered only upon total equipment failure or violation of safety limits, making them incapable of signaling gradual degradation of energy performance.

To fill the gap in hindsight, energy management systems (EMS) have emerged in response to the growing need to visualize data collected by smart meters. These platforms shift their focus to data aggregation, periodic reporting, and operational benchmarking. However, Zhao et al. [5] highlight a significant limitation of these systems: their inherently descriptive nature. While an EMS can successfully indicate that energy consumption has increased on a given day, it lacks the logical mechanisms necessary to isolate the underlying technical cause, thus leaving the entire investigation task to the human operator.

The next level of analytical complexity within this ecosystem is fault detection and diagnosis (FDD) solutions. The

industry standard approach to implementing FDD remains highly deterministic and rule-based, using extensive libraries of logical expressions. Although these systems are well validated for detecting obvious mechanical faults, their widespread application faces major obstacles, mainly due to the high costs of initial setup and ongoing maintenance. Furthermore, as Granderson et al. [4] have noted, static rules quickly become obsolete as buildings undergo natural functional changes over time. Moreover, these rigid logic structures fail to integrate complex contextual variables, frequently leading to false positives during atypical events or extreme weather conditions.

### *C. Transition to Data-Driven Analysis*

The inherent limitations of deterministic systems have led to a shift in research efforts towards data-driven methodologies. This shift is significantly supported by the widespread availability of advanced measurement infrastructure (AMI), which generates high-resolution time series. Unlike physical models that require hard-to-obtain parameters, these data-driven approaches use numerical history to automatically extract and analyze behavioral patterns.

Historically, a major obstacle to the validation of these modern energy efficiency algorithms has been the lack of standardized and open datasets. Therefore, for an anomaly detection algorithm to be rigorously and scientifically validated today, it is imperative that it be tested on a wide range of standardized consumption profiles [6]. The Building Data Genome 2 (BDG2) dataset [7] represents a comprehensive and representative resource for the building energy analytics community which aggregates hourly data from over 3,000 energy meters. The main importance of this dataset lies in the heterogeneity and realistic quality of the data, which inherently contain real-world noise and operational gaps. Using the BDG2 dataset ensures that any proposed solution possesses the generalizability required for real-world, large-scale applications.

Finally, the large volume and variability of the BDG2 data clearly demonstrated that simple statistical methods are insufficient for accurate analysis. Commercial buildings exhibit highly nonlinear behavior that is continuously influenced by multiple competing factors. Therefore, as studies derived from this dataset support [7], [8], machine learning algorithms are uniquely capable of modeling this inherent complexity and learning subtle operational correlations without the need for explicit programming.

### *D. Unsupervised Learning Algorithms and the Matrix Profile*

In an unsupervised learning scenario, anomaly detection fundamentally involves identifying subsequences that deviate significantly from the typical operating pattern. However, traditional clustering methods, such as K-Means, have proven to be extremely inefficient for this specific task, as they tend to “average” temporal behavior, thereby masking rare but critical events [9]. Furthermore, while deep neural networks offer advanced capabilities, they often operate as “black boxes,” providing results that lack the necessary transparency and are

therefore difficult for facility managers to interpret and act upon.

To directly address these analytical limitations, Yeh et al. [10] introduced Matrix Profile, a data structure that annotates a time series with the exact distance of its nearest neighbor. Unlike traditional heuristic methods, Matrix Profile provides a strict mathematical guarantee for discovering the most similar and most different sequences in a data set. This robust approach elegantly solves two fundamental problems in time series analysis: identifying “patterns,” which represent normal recurring behavior, and detecting “discrepancies,” which are unique subsequences that indicate potential operational anomalies. Specifically, this research uses the STUMPY library [11], a modern implementation of Matrix Profile using the Python programming language that accelerates the use and applicability of the algorithm.

### III. METHODOLOGY AND ALGORITHMS

To achieve the research objectives, a structured workflow was adopted, consisting of four main stages: data selection, preprocessing, algorithmic analysis, and integration into a user interface. This methodological framework ensures the transition from raw historical data to validated and actionable information.

#### A. Dataset Selection

For the experimental validation of the proposed system, the building identified as “Hog\_office\_Byron” from the Building Data Genome 2 (BDG2) dataset was selected as reference model. The selection criteria focused on the specific characteristics of the building to ensure relevance:

- 1) Use Type: Commercial office building, characterized by distinct weekly consumption patterns.
- 2) Climate Zone: Located in a temperate climate (Zone 6A – Cold and Wet), which implies significant seasonal variations in energy demand for both heating and cooling.
- 3) Data Availability: The dataset covers a 24-month period with hourly resolution, providing sufficient historical depth for the algorithm to learn long-term seasonality.

#### B. Exploratory Data Analysis

Before applying the advanced detection algorithms, an initial visual inspection was performed using a heatmap (Figure 1). This technique maps the time of day to the day of the year, assigning a color intensity proportional to energy consumption.

The analysis revealed a clear operational schedule (08:00–18:00), consistent with office activities. In addition, the visualization highlighted a significant data quality issue: a prominent horizontal band of missing values spanning a continuous period of several weeks, which required subsequent data imputation.

#### C. Anomaly Detection with Matrix Profile

The main analytical component of the system uses the Matrix Profile algorithm, implemented through the STUMPY library. The technical basis of this implementation is based

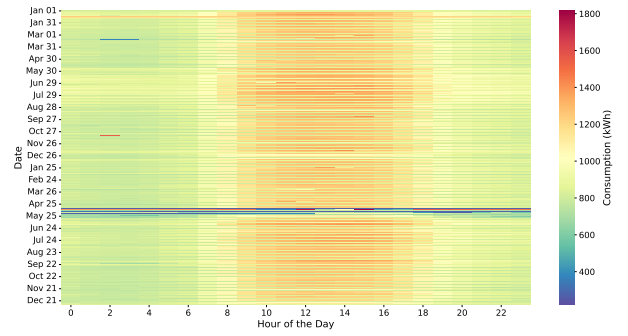


Fig. 1. Energy consumption heatmap for the BDG2 Hog\_office\_Byron building.

on the use of the Z-normalized Euclidean distance between subsequences, a critical step that allows the algorithm to identify shape similarities (motifs) and discrepancies (anomalies) regardless of amplitude variations or linear delays in energy consumption. The Euclidean distance between two subsequences  $T_a$  and  $T_b$  of length  $n$  is defined as:

$$d(T_a, T_b) = \sqrt{\sum_{i=1}^n (T_{a,i} - T_{b,i})^2} \quad (1)$$

For this study, a window length,  $n$ , of 24 hours was selected to capture the daily diurnal cycle. To ensure mathematical accuracy and prevent “trivial matches” - where the algorithm identifies an adjacent overlapping subsequence as a neighbor - a temporal exclusion zone of  $m/2$  was implemented. The processing generates two primary data structures: the Matrix Profile vector, which stores the calculated minimum distances, and the Matrix Profile index, which maps the index of the nearest neighbor for each analyzed temporal window. This unsupervised framework provides a robust method for discovering operational irregularities without the need for pre-trained models or manually labeled historical datasets [12].

The results of applying the algorithm are visible in Figure 2, where the anomaly score (bottom graph) highlights the critical moments.




Fig. 2. Result of applying the Matrix Profile algorithm.

#### D. Climatic Analysis (Energy Signature)

To distinguish between mechanically induced anomalies and weather-driven consumption peaks, the methodology integrates an analysis of the building's energy signature. By correlating daily energy averages with external temperature data, a scatter plot is generated.




Fig. 3. Building energy signature.

This visualization typically reveals a double dependency in temperate climates: consumption increases as temperatures fall below a heating equilibrium point (approximately  $10^{\circ}\text{C}$ ) and increases again as temperatures exceed a cooling threshold ( $20^{\circ}\text{C}$ ). Deviations from the expected operational reference level serve as a secondary layer of validation for alerts generated by the Matrix Profile algorithm.

#### IV. SOFTWARE IMPLEMENTATION

The realization of the proposed system required the adoption of a modular software architecture, capable of handling the complex data flow, from the ingestion of raw time series to the visualization of performance indicators. The system was developed using the Python 3.14 programming language, selected for its extensive ecosystem of libraries dedicated to data science and numerical computing [13].

The operational link between the mathematical foundation and the user interface is achieved through the Decision Engine. The backend processing pipeline serializes the distance vector from the Matrix Profile and passes it to the decision engine, which applies a dynamic threshold, calculated as a multiple of the average of the distances. When a value in the Matrix Profile exceeds this threshold, the system not only signals an anomaly, but also automatically triggers the calculation scripts for financial and ESG impact, transforming an abstract mathematical distance into a quantifiable maintenance ticket.

##### A. Backend Architecture and Design

Figure 4 illustrates the high-level software architecture of the proposed decision support system. The system is designed based on a layered framework that orchestrates the transformation of raw building data into concrete facilities management information.

The data flow starts at the data source level, which ingests heterogeneous time series from the Building Data Genome Project 2 (BDG2) and external meteorological records. These

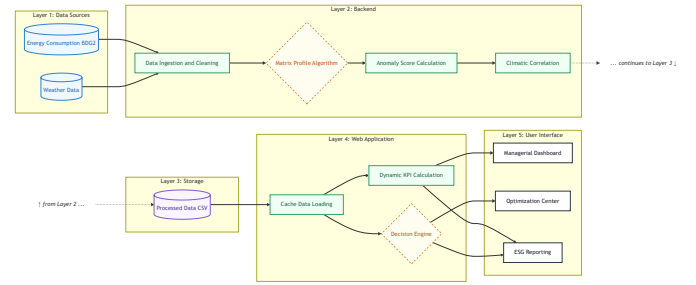


Fig. 4. High-level software architecture of the developed decision support system.

inputs are processed at the backend, where an Extract-Transform-Load (ETL) pipeline cleans the data before sending it to the Matrix Profile algorithm, implemented through the stumpy library - to calculate anomaly scores and establish climatic correlations. To ensure system responsiveness, the processed outputs are serialized to the storage layer, decoupling computationally intensive tasks from user interactions. The web application layer, developed using the Streamlit framework, hosts the main decision engine and dynamic KPI calculations e.g. EUI, financial impact, carbon footprint. Finally, the user interface layer renders these analyses through an interactive dashboard [14].

To facilitate visual tracking of the information flow and highlight the separation of concerns, the architectural design uses a semantic color and shape coding scheme. Specifically, blue nodes represent raw data sources, such as CSV files, while green nodes denote computational processing tasks, including data cleaning, anomaly scoring, and dynamic KPI calculations. The algorithmic core of the system and decision logic are explicitly highlighted by orange dotted diamonds, which indicate complex conditional or mathematical evaluation points. Finally, the purple components indicate the data persistence layers used for storage efficiency, and the white elements describe the user interface modules through which the plant operator interacts with the system.

##### B. Backend Processing

The processing component operates as a batch process, running independently of the graphical user interface (GUI). The implemented data pipeline includes the following distinct steps:

- 1) *Ingestion and Cleansing (ETL)*: Python scripts use the Pandas library to load datasets from CSV format. In this step, the consumption time series is aligned with the weather data, and missing values are treated by linear interpolation to ensure continuity of the analysis algorithms.
- 2) *Matrix Profile Calculation*: For each 24-hour time window, the system calculates the Euclidean distance to the historical nearest neighbor. The result is a vector of anomaly scores, where high values indicate discord sequences.

- 3) *Climate Correlation*: The backend integrates external temperature data to contextualize consumption. Through timestamp-based connection operations, each consumption record is associated with the corresponding temperature, allowing for additional filtering of alerts caused by extreme weather conditions..

### C. Interactive Frontend

The user interface was built using the Streamlit framework [14], a modern solution that allows for the rapid transformation of data scripts into interactive web applications. Designed with a focus on ease of use and clarity, the platform targets the specific operational needs of a facilities manager. For advanced data graphing, the Plotly library is integrated, providing interactive functionality superior to static charts, allowing the user to zoom in, debug, and inspect individual values through hover actions. In addition, a significant aspect of the implementation is the adaptability of the interface to the specific theme. By using native CSS variables, the visual style automatically adapts to the user’s system preferences, switching between dark and light modes to minimize visual fatigue for operators who monitor screens for long periods of time.

In addition to its visual components, the interface extends beyond static data display by incorporating dynamic computational logic. The system features a real-time calculation engine that instantly recalculates key performance indicators, such as EUI, as well as financial impact estimates. These dynamic updates are triggered directly by structural and operational parameters entered by the user in the sidebar, such as gross built-up area and facade type, ensuring a highly responsive analytical environment.

## V. EXPERIMENTAL RESULTS

The developed application was validated using the reference historical data set from the “Hog\_office\_Byron” building as described beforehand.

To validate the benchmarking capability, an algorithm was developed to select similar buildings that minimize the difference in floor area within the same use category. Thus, the reference building *Hog\_office\_Byron* (80,038  $m^2$ ) was compared with the building *Rat\_office\_Colby* (79,000  $m^2$ ). Although the difference in floor area is negligible (1.3%), the analysis revealed a major performance discrepancy: the similar building has an energy consumption of 472.7  $kWh/m^2$ , 126.1% higher than the reference (209  $kWh/m^2$ ). This comparison, supported by an anomaly score of 4.30 for the similar building, demonstrates that structural similarity does not guarantee operational efficiency, justifying the need for automated diagnostic tools to isolate massive energy waste (equivalent to € 3,091,568/year).

### A. Main Dashboard and KPI Monitoring

The central module (Figure 5) provides a complete overview of the building’s operational status. Unlike traditional BMS screens that typically display instantaneous values such as

current voltage or air flow, this dashboard aggregates data to calculate dynamic key performance indicators (KPIs). For example, the system calculates an energy use intensity (EUI) of 327  $kWh/m^2$  for the analyzed period. This specific measurement is dynamically normalized based on the gross floor area (GFA) entered by the user, facilitating immediate comparison with established industry standards.

In addition, the module incorporates an automatic cost estimation function; by applying a standard commercial rate, it converts gross kilowatt-hours directly into monetary value to show the operating expenses since the beginning of the year. Finally, to address the common absence of physical sub-metering, the system generates a consumer fault. It estimates the distribution of energy consumption in categories such as HVAC, lighting and outlet loads using probabilistic models adapted to the specific building type and time of day.

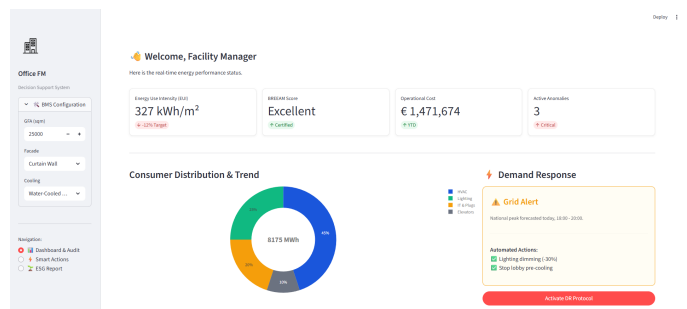


Fig. 5. Main system dashboard.

The EUI calculation (327  $kWh/m^2$ ), the estimated consumption breakdown and the demand response module are visible.

### B. Optimization Center (Decision Support)

This module represents the main innovation of the proposed system. Instead of relying on passive graphs, the application interprets the Matrix Profile anomaly scores to generate actionable tickets. During the testing phase, the system successfully identified two distinct types of inefficiencies and automatically quantified their financial impact using specific cost modeling algorithms.

*Soft Start Opportunity*: The algorithm detected a recurring increase in consumption at 05:00 during the winter months. The decision engine marked this event as a “Hard Start” event and generated a recommendation to implement a “Soft Start” sequence. The system estimated a potential saving of approximately 29,800€/year. This estimate is derived from the peak tariff reduction, calculated using the following equation, where  $S_{peak}$  represents the potential financial savings:

$$S_{peak} = P \times C \quad (2)$$

Here, the variable  $P$  indicates the power surplus, defined as the difference in kilowatts between the recorded peak and the nominal load. The variable  $C$  represents the penalty cost, which is the surcharge applied by the utility provider for exceeding the subscribed power limit.

**Phantom Load Detection:** The analysis revealed a consistently high consumption during weekends (Saturday-Sunday), deviating from the expected reference load. The system classified this as “Phantom Load” and recommended a physical audit of the local HVAC units on the upper floors. To motivate the intervention, the financial loss is calculated by isolating the non-operational energy waste using the formula below, where  $S_{phantom}$  is the estimated savings:

$$S_{phantom} = E \times T \quad (3)$$

In this equation,  $E$  corresponds to the weekend consumption, isolating the non-operational energy waste measured in raw energy data (kWh), and  $T$  is the current energy tariff used to convert this data directly into a monetary value. This application of this formula in our case yields a potential waste reduction of 15%.

### C. ESG and Sustainability Reporting

To align with modern corporate reporting requirements, the system includes a dedicated module for Environmental, Social and Governance (ESG) criteria. The application automatically converts energy consumption into carbon emissions (tonnes of CO<sub>2</sub>) using a conversion factor specific to the local energy mix (set to 0.25 kg/kWh for this experiment), from [15].

To make the data relatable to non-technical stakeholders, the interface visualizes the environmental impact in equivalent terms. As illustrated in Figure 6, the module displays the tangible equivalence between the total emissions of the building and the number of trees needed to sequester the emitted carbon. This feature directly facilitates communication between facility managers and boards of directors on sustainability goals and reporting.




Fig. 6. ESG reporting module.

## VI. CONCLUSIONS AND FUTURE DIRECTIONS

The research validated the feasibility of integrating the Matrix Profile algorithm into a decision support system (DSS) capable of identifying operational anomalies without the need for labeled data sets. Testing on a reference building demonstrated that the use of open-source technologies allows the transformation of raw time series into practical maintenance actions. The implementation generated immediate quantitative results, identifying potential savings of €29,800/year through

equipment start-up optimization and a 15% reduction in “Phantom Load” waste. At the same time, the integration of the 0.25 kg/kWh conversion factor allowed the monitoring of a carbon footprint of 2,044 tonnes of CO<sub>2</sub>, providing a rigorous basis for ESG reporting.

The main limitation of the current prototype is the offline architecture and the lack of sub-metering, which requires the use of probabilistic models for energy usage. Future research directions will focus on the transition from descriptive to predictive analytics by implementing forecasting models, such as LSTM [16] or Prophet, and direct connection to BMS systems via industrial protocols such as BACnet or MQTT. Continuous validation through “Human-in-the-loop” mechanisms will ensure algorithm refinement and reduction of false alarm rates in real operating conditions.

## REFERENCES

- [1] G. Cretu, I. Stamatescu, and G. Stamatescu, “Modeling and prediction of occupancy in buildings based on sensor data using deep learning methods,” *IEEE Access*, vol. 12, pp. 102 994–103 003, 2024.
- [2] S. Katipamula and M. R. Brambley, “Methods for fault detection, diagnostics, and prognostics for building systems—a review, part i,” *Hvac&R Research*, vol. 11, no. 1, pp. 3–25, 2005.
- [3] M.-G. Boicu, G. Stamatescu, I. Făgărășan, M. Vasluianu, G. Neculoiu, and M.-A. Dobrea, “Optimizing occupant comfort in a room using the predictive control model as a thermal control strategy,” *Sensors*, vol. 24, no. 12, 2024. [Online]. Available: <https://www.mdpi.com/1424-8220/24/12/3857>
- [4] J. Granderson and G. Lin, “Building energy information systems: synthesis of costs, savings, and best-practice uses,” *Energy Efficiency*, vol. 9, no. 6, pp. 1369–1384, 2016.
- [5] H.-x. Zhao and F. Magoulès, “A review on the prediction of building energy consumption,” *Renewable and Sustainable Energy Reviews*, vol. 16, no. 6, pp. 3586–3592, 2012.
- [6] M. Bourdeau, X. qiang Zhai, E. Nefzaoui, X. Guo, and P. Chatellier, “Modeling and forecasting building energy consumption: A review of data-driven techniques,” *Sustainable Cities and Society*, vol. 48, p. 101533, 2019.
- [7] C. Miller, A. Kathirgamanathan, B. Picchetti, P. Arjunan, J. Y. Park, Z. Nagy, P. Raftery, B. W. Hobson, Z. Shi, and F. Meggers, “The building data genome project 2, energy meter data from the ashrae great energy predictor iii competition,” *Scientific data*, vol. 7, no. 1, p. 368, 2020.
- [8] D. B. Araya, K. Grolinger, H. F. ElYamany, M. A. Capretz, and G. Bitsuamlak, “An ensemble learning framework for anomaly detection in building energy consumption,” *Energy and Buildings*, vol. 144, pp. 191–206, 2017.
- [9] E. Keogh and J. Lin, “Clustering of time-series subsequences is meaningless: implications for previous and future research,” *Knowledge and information systems*, vol. 8, no. 2, pp. 154–177, 2005.
- [10] C.-C. M. Yeh, Y. Zhu, L. Ulanova, N. Begum, Y. Ding, H. A. Dau, D. F. Silva, A. Mueen, and E. Keogh, “Matrix profile i: all pairs similarity joins for time series: a unifying view that includes motifs, discords and shapelets,” in *2016 IEEE 16th international conference on data mining (ICDM)*. Ieee, 2016, pp. 1317–1322.
- [11] S. Law, “STUMPY: A Powerful and Scalable Python Library for Time Series Data Mining,” *The Journal of Open Source Software*, vol. 4, no. 39, p. 1504, Jul. 2019.
- [12] G. Stamatescu, R. Plamanescu, I. Ciornei, and M. Albu, “Detection of anomalies in power profiles using data analytics,” in *2022 IEEE 12th International Workshop on Applied Measurements for Power Systems (AMPS)*, 2022, pp. 1–6.
- [13] W. McKinney *et al.*, “Data structures for statistical computing in python,” *scipy*, vol. 445, no. 1, pp. 51–56, 2010.
- [14] A. Treuille, T. Teixeira, and A. Kelly, “Streamlit: The fastest way to build and share data apps,” *Streamlit Documentation*, 2020.
- [15] IPCC, “Gwp 100a emission factors,” IPCC, Tech. Rep., 2013.
- [16] S. Hochreiter and J. Schmidhuber, “Long short-term memory,” *Neural computation*, vol. 9, no. 8, pp. 1735–1780, 1997.