Data and Knowledge–Driven Approach for Energy Profiling in Smart Context-Aware Buildings

Mona Farrag[1*], Gerald Feldman[1], Haitham Mahmoud[1], Nouh Elmitwally[1], and Mohamed M. Gaber[2]

[1] College of Computing, Birmingham City University, Birmingham, United Kingdom
mona.farrag@mail.bcu.ac.uk, gerald.feldman@bcu.ac.uk
[2] Queensland University of Technology, Brisbane, Australia

Abstract. Energy profiling plays a crucial role in optimising smart building operations, especially with the increasing popularity of personalised, user-centric AI applications. Current research lacks emphasis on interpretability, transparency, and accessibility for non-expert stakeholders, where decision-making either relies solely on machine learning insights or unstructured knowledge bases. Hence, this study aims to enhance the interpretability of energy profiling and generate tailored recommendations based on correlated data sources from various aspects. This approach combines data-driven and knowledge-driven techniques by integrating energy clustering insights and unstructured knowledge bases to provide tailored energy recommendations. By combining Large Language Models (LLMs) and Explainable AI (XAI), this approach leads to: (1) identifying new consumer personas based on contextualised cluster insights, (2) finding the most impactful features reflecting energy insights, and (3) turning those insights into clear, human-readable reports and recommendations. This transforms smart meters from passive data collectors into intelligent advisory tools for consumers, policymakers, and energy providers.

Keywords: Smart building · Energy profiling · Smart Meter Analysis · Large Language Models · Explainable AI.

1 Introduction

The evolution of smart building technologies, particularly the widespread deployment of smart meters, has introduced new opportunities for optimising energy consumption at the household level. As of September 2024, over 37 million smart and advanced meters had been installed in homes and small businesses across Great Britain, representing approximately 65% of all meters in use [8]. These smart meters generate continuous time-series data, capturing changing electricity usage patterns and consumption trends at set intervals, and transmit real-time data to utilities [18]. Moreover, studies have shown a strong link between smart meter feedback and occupant behaviour [3]. Many households, for instance, willingly adjust their electricity usage during peak hours once they understand their consumption patterns, indicating a clear potential for behavioural adaptation based on data transparency.

However, despite the growing volume and resolution of smart meter data, significant challenges persist in interpreting this information in a meaningful way. Current research often isolates machine learning models from contextual knowledge or natural language explanations. Data-driven approaches provide accurate pattern detection but tend to fall short in interpretability and accessibility for non-expert users. Similarly, LLMs and natural language processing (NLP) techniques have been applied independently for summarisation or recommendation tasks, but they struggle with numeric reliability, lack of context grounding, and inconsistent outputs when used in isolation. As a result, existing tools frequently fail to convert raw data into actionable insights that can support energy-saving decisions for consumers or inform policymaking. To complement this, XAI can bridge the gap between model outputs and human understanding. Hence, this study proposes a holistic, AI-driven hybrid pipeline that combines XAI with LLMs to overcome the non-determinism of LLMs and enhance interpretability and human readability for machine learning models. Providing contextual insights into the typical consumers and their personas, along with tailored energy-saving recommendations. Incorporating explainable AI to create a data- and knowledge-driven analysis framework that promotes transparency and insight generation across multiple dimensions.

2 Related Works

2.1 Behavioural Energy Profiling and Clustering Techniques

Machine learning and data-driven approaches play a pivotal role in residential energy profiling by identifying distinctive consumption patterns and capturing occupant behaviours. For instance, the work presented in [12] employed Support Vector Regression (SVR) to forecast energy usage across various temporal and spatial resolutions. While this approach effectively captured nuanced consumption trends, it is constrained by its dependency on historical data, high computational demands, and limited generalisability to novel or irregular data scenarios. To address the challenge of personalising energy interventions, [14] proposed a method for recommending electricity tariffs based on demographic and behavioural clustering using K-means and decision trees. This model enabled actionable recommendations even in the absence of historical smart meter data. However, the approach is sensitive to the quality and accuracy of demographic input data, and its practical deployment is hampered by difficulties in validating outputs due to the lack of historical benchmarks. In parallel, qualitative investigations have offered valuable insights into human-centric aspects of energy consumption. For example, [1] adopted a human-centred methodology incorporating surveys and interviews to develop occupant personas that integrate physical, physiological, and psychological traits. These personas have informed smart housing design by embedding contextual understanding into system-level decisions. Nevertheless, such qualitative methods often suffer from scalability limitations, subjectivity biases, and challenges in translating qualitative insights into actionable metrics for large-scale deployment. A nuanced understanding of occupant behaviour remains crucial for optimising building energy performance. Various machine learning models have been explored to characterise such behaviours through clustering techniques. For example, the study in [15] utilised unsupervised learning on U.S. residential data to cluster consumption patterns. While the approach was effective in identifying usage archetypes, it lacked explainability mechanisms, limiting interpretability and trust in the decision-making process.

2.2 Explainability and Language Models in Energy Forecasting

In [20], model explanations are closely tied to the convolutional neural network (CNN) architecture, limiting transferability across models. The study in [2] combines NLP with ExtraTrees regression and Local Interpretable Model-agnostic Explanations (LIME) for the prediction of electricity demand on a day-to-day basis, but is limited by the reliance on high-quality textual data, computational demands, and limited causal interpretability. Similarly, [23] integrates CNN and Random Forest models with SHapley Additive exPlanations (SHAP) and LIME to provide electricity consumption feedback. However, the lack of occupant context, such as routines and lifestyle, reduces the relevance and specificity of recommendations. LLMs have also been explored for forecasting tasks. In [22], LLMs are used for electricity prediction, while [13] applies them to human activity classification. These studies aim to leverage LLMs' contextual reasoning capabilities, but face issues such as numeric instability, prompt sensitivity, and poor performance with outliers. Small prompt variations can lead to inconsistent outputs, making LLMs unreliable for precise numerical forecasting and applications requiring stable long-term historical analysis.

3 Methodology

This study presents a holistic approach that integrates data-driven clustering, explainable AI, and LLMs to analyse household energy consumption patterns. The process begins with data exploration and cleaning, followed by time-series aggregation and clustering of households based on daily energy usage using K-Means. Environmental factors such as weather and seasonality are then incorporated to enhance contextual understanding. SHAP analysis is applied to interpret feature importance within each cluster. Finally, the outputs of the graphs and insights are processed by two LLMs to generate descriptive reports and tailored recommendations, combining data analysis with domain knowledge to achieve actionable outcomes, as shown in Figure 1.

3.1 Data Sources

This study utilises three data sources: (1) Smart meter data $\mathcal{D}_{\text{smart}}$ from 5,567 London households, focused on 4,040 standard tariff users during 2013, denoted as: $\mathcal{D}_{\text{smart}} = \{x_i^t \in \mathbb{R} \mid i = 1, \ldots, n, \;\; t = 1, \ldots, T\}$ where x_i^t is the half-hourly energy consumption of household i at time t. (approximately 167 million rows in CSV format) [5] ; (2) Weather data $\mathcal{D}_{\text{weather}}$ containing daily variables: $\mathcal{D}_{\text{weather}} = \{w^t = [\text{temp}_t, \text{radiation}_t, \text{cloud}_t] \in \mathbb{R}^3 \mid t = 1, \ldots, T'\}$ (resulting in 365 rows in the sliced data frame) [19]; and (3) A knowledge base $\mathcal{K}$ from energy providers with domain-specific practices and policies [7]. It contains information on load balancing, the environmental benefits of shifting power use to off-peak hours, and practical suggestions aligned with real-world applications.

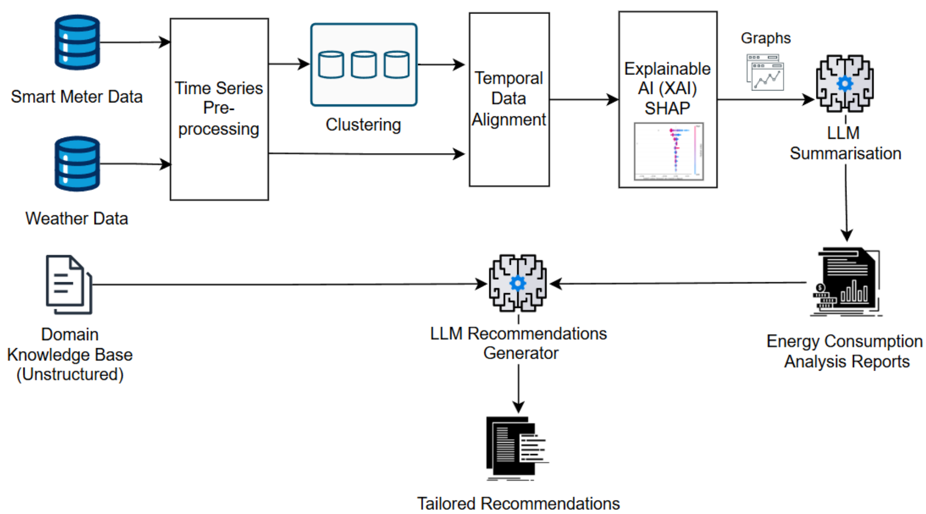

Fig. 1: AI-driven energy consumption analysis and recommendation pipeline.

3.2 Data Driven Time Series Analysis

Building upon our methodological framework, Figure 2 elaborates on the data-driven analytical components within the broader pipeline shown in Figure 1, showing the sequence of operations performed from initial data preparation through clustering and explainability analysis.

- **Data Exploration and Cleaning**: The data quality is assessed by exploring missing values, identifying relevant subsets, and filtering the data. The data frame is sliced for the specific time scope of interest. Missing values in both datasets are identified, and smart meters with more than a threshold

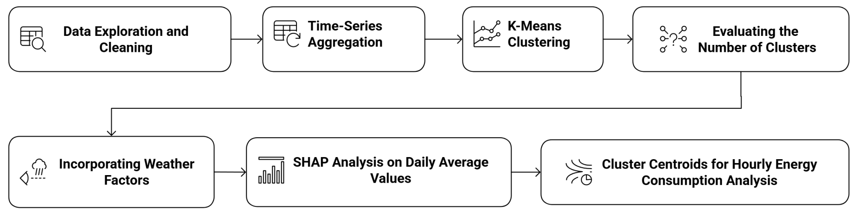

Fig. 2: Methodological workflow for the data driven smart meter time series consumption analysis.

of 3 missing values are eliminated using the proposed rule, ensuring that households with excessive missing data are excluded from the dataset. The remaining missing values are imputed using mean values. Households are filtered based on the standard tariff. Missing values $\text{MV}(x_i)$ are handled by:

$$x_i^t = \begin{cases} \text{mean}(x_i), & \text{if } x_i^t \text{ is missing and } \#\text{MV}(x_i) \leq 3 \\ \text{discard}, & \text{if } \#\text{MV}(x_i) > 3 \end{cases}$$

- **Time-Series Aggregation**: Since the dataset contains half-hourly smart meter readings, the data is down-sampled to create daily data for long-term consumption trends and environmental impact analysis [9]. Using hourly data to study intraday consumption behaviour within weekdays and weekends in different seasons. Half-hourly data is aggregated to daily consumption: $X_i^d = \sum_{h=1}^{48} x_i^{(d-1)\cdot 48+h}, \quad d = 1, \ldots, D$ where X_i^d is the daily energy consumption of household i on day d.
- **Clustering Households Based on Energy Usage for Daily Consumption Over a Year**: To segment households with similar energy usage patterns, K-means clustering is applied. K-means clustering has been widely used in modern power systems for tasks such as load forecasting, fault detection, power quality analysis, and system security assessment. It is an efficient and scalable clustering technique, chosen for its reliability in time-series analysis. SHAP selects features that ensure better model accuracy [21]. Its shape-preserving capabilities ensure that clustered consumption patterns retain their characteristic fluctuations, making it particularly effective for capturing temporal resolution effects in energy usage data. The energy data is standardised using z-score normalisation. K-means clustering is then applied to group similar households. The optimal number of clusters is determined using the elbow method, which suggests that the appropriate number of clusters falls within the range of 4 to 6. For evaluation, 6 clusters are selected. Let $\mathbf{X} \in \mathbb{R}^{n \times D}$ be the aggregated data matrix. K-Means minimises: $\min_{C_1,\ldots,C_k} \sum_{j=1}^{k} \sum_{x_i \in C_j} \|x_i - \mu_j\|^2$ where μ_j is the centroid of cluster C_j. The optimal k is selected via the elbow method.
- **Incorporating Environmental Factors**: To assess the significance of the weather influences energy consumption, the daily weather dataset is merged

with the clustered energy data. Daily weather features, including temperature, radiation, and cloud cover, are combined with energy data. Seasonal variations are analysed by adding a "Season" column, and bank holidays and weekends are identified for further analysis. Each data point is appended with weather features: $Z_i^d = [X_i^d, \text{temp}_d, \text{radiation}_d, \text{cloud}_d]$

- **Explainability for Each Cluster with SHAP Analysis**: SHAP is a model explainability technique that assigns an importance score to each feature, encoding the impact of that feature on the model [17]. In this study, clustering results are transformed into classification tasks to enable the effective application of explainability techniques, following the approach of [4]. For each cluster, a classifier $f(Z) \to y$ is trained and SHAP assigns contributions ϕ_j to features: $f(Z_i) = \phi_0 + \sum_{j=1}^{d} \phi_j(Z_i)$ where ϕ_0 is the base value. This transformation allows the mapping of environmental weather features to the clusters, enabling interpretation of the external conditions associated with different usage patterns, forming a Multivariate Time Series (MTS). MTS is a type of time series where multiple variables (features) are recorded simultaneously over time. Inter-variable relationships refer to the correlation and influence between variables over time[10]. SHAP is then applied to each cluster to understand the importance of each feature in the MTS and to clarify inter-variable relationships.
- **Hourly Energy Consumption for Day Type and Bank Holidays**: To assess weekend vs. weekday patterns, hourly data is grouped by weekdays, weekends, and bank holidays. In addition, the monthly trend is used to identify changes in energy demand across different times of the year. The graphs generated from this analysis are passed to the LLM unit to produce interpretable summaries. Let x_i^h be the hourly consumption. We group by day type: $\text{MeanUsage}_{\text{type},h} = \frac{1}{|\mathcal{H}_{\text{type}}|} \sum_{i \in \mathcal{H}_{\text{type}}} x_i^h$ where $\mathcal{H}_{\text{type}}$ denotes the set of households for weekdays, weekends, or holidays.
- **Centroid Computation:** The centroid for each cluster has been calculated for a given cluster $\mathcal{C}_k$ using the formula: $\text{Centroid}_k(t) = \frac{1}{N_k} \sum_{i \in \mathcal{C}_k} \text{Energy}_{i,t}$ Where $\text{Centroid}_k(t)$ is the centroid of cluster k at time t. N_k is the number of households in cluster k. $\text{Energy}_{i,t}$ is the energy consumption of household i at time t.

3.3 Evaluating Clusters

The Silhouette Score plot was utilised to determine the optimal number of clusters. The silhouette analysis revealed that $k = 2$ achieved the highest silhouette score of 0.5210, indicating that two clusters may represent the most optimal configuration in terms of well-defined separation. However, the silhouette score exhibited a steady decline as k increased, indicating that the introduction of additional clusters led to less distinct groupings. . Although $k = 6$ produced a silhouette score of 0.3038 higher than the values observed for $k \geq 7$, it was not the mathematically optimal choice, as $k = 2$ and $k = 3$ yielded slightly higher scores. Nevertheless, this study prioritises interpretability and exploratory analysis over strict optimisation. Following the approach of [6], where researchers did not rely solely on the silhouette's optimal cluster number but instead selected a larger k to uncover additional insights, we likewise chose to proceed with 6 clusters. We compute the Silhouette Score $s(k)$ for each k value: $s(k) = \frac{b(k)-a(k)}{\max(a(k),b(k))}$, where $a(k)$ is the average intra-cluster distance and $b(k)$ is the minimum average inter-cluster distance [16].

3.4 LLMs for Knowledge-Driven Insights

The generated analysis and SHAP for each cluster are initially exported in PDF format. These documents serve as input to the LLM unit, which utilises GPT-4o to produce descriptive summaries for the graphical representations associated with each cluster. While XAI contributes interpretability, it often lacks contextual reasoning and user-centred narrative delivery; integrating LLMs addresses these limitations, as supported by [11]. Following the image summarisation process, the explained graphs, along with their corresponding textual summaries, are compiled into a PDF report for each cluster. These reports, which integrate insights derived from XAI techniques and the generative language model, are further refined using GPT-3.5-turbo. Let G_j be the graph summaries and T_j be their textual descriptions for cluster j. GPT-4o generates: $S_j = \text{GPT-4o}(G_j) \quad \Rightarrow \quad R_j = \text{GPT-3.5-turbo}(S_j, \mathcal{K})$ where S_j is a summarised report and R_j is the refined recommendation. This produces interpretable, actionable advice aligned with contextual behaviours derived from data and knowledge integration.

Through carefully designed prompt engineering, this phase combines the analytical findings with actionable recommendations sourced from a domain-specific textual knowledge repository, also stored in PDF format. This integration enhances the relevance and applicability of the generated recommendations. The process of developing an explained analysis with recommendations for the created energy profiles which is illustrated in Figure 3, where prompt engineering operates through two integrated phases:

- **Prompt Engineering for Explainer Unit**: To integrate LLMs into the analysis, we employed a two-part prompt design. The system prompt defined the model's role as "a bot that is good at analysing graphs and images for smart meter time series data, linking them to contextual text for energy profile analysis". The user prompt combined a generic instruction "Analyse the graph and extract all insights from it?" with the input analysis graphs and SHAP figures. In this step, figures and structured prompts were combined to generate analytical insights, which were then stored as structured summaries for the recommendations function.
- **Prompt Engineering for Advisor Unit**: Building on the previous prompt, a second prompt structure was implemented to produce household-level recommendations. The system prompt defined the model as "an energy assistant that provides clear, actionable, and context-aware energy saving recommendations for households". The user prompt combined the textualised insights generated from the analysis stage (image summaries and SHAP explanations

converted to text) with content extracted from a curated knowledge-base PDF of energy efficiency guidance (cleaned and converted to text). This was operationalised through a function `generate_recommendations`(I, K), which dynamically inserts both the textualised insights I and the knowledge-base text K into a fixed prompt template, thereby generating tailored insights for each energy cluster.

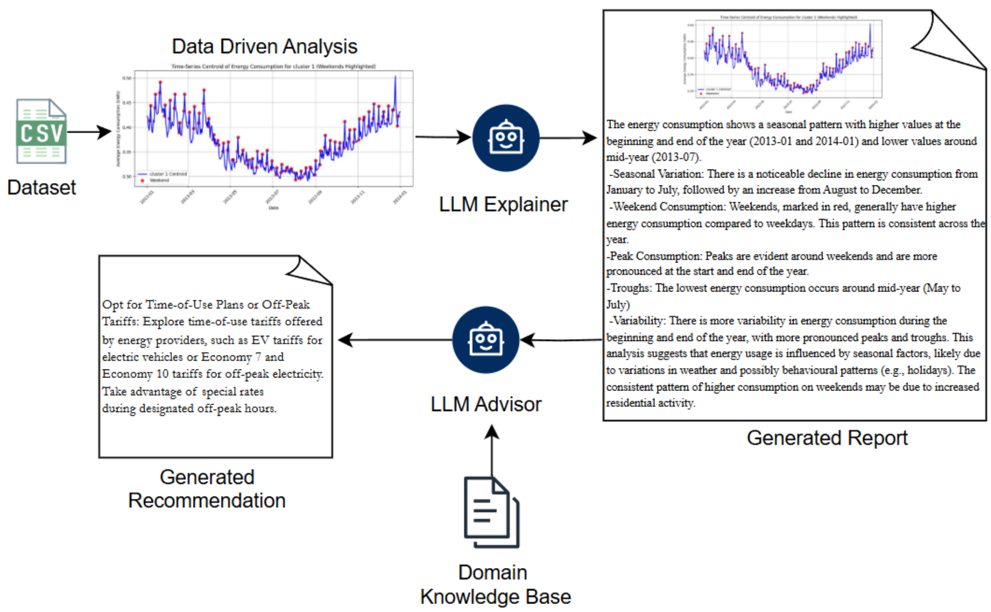

Fig. 3: Illustration of the process of generating explained energy consumption analysis and recommendations using LLMs.

4 Results and Discussion

4.1 Results

This approach allows the explainable model to uncover hidden patterns and thoroughly test the proposed methodology. The analysis reveals clear temporal patterns in energy consumption, supporting the claim that SHAP and LLMs significantly enhance the transparency and interpretability of the analysis, highlighting insights that may be overlooked using traditional methods. The centroid graphs for the clusters, shown in Figure 4 further illustrate seasonal trends in energy consumption.

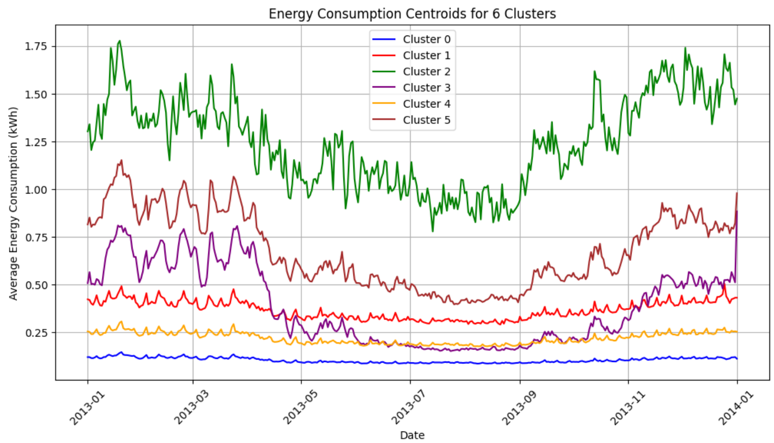

Fig. 4: Energy consumption centroids for the six clusters of the energy profiles over time

In the presented graph, Cluster 2 shows the highest energy usage, with consumption peaking in winter and dropping in summer, likely due to heating needs. Cluster 5 follows a similar pattern but with slightly lower consumption. Clusters 3 and 1 exhibit moderate changes, indicating routine-based energy use, while Clusters 4 and 0 have the lowest consumption, with slight seasonal variation, suggesting smaller households or more energy-efficient habits. Based on this study, selecting a higher K-value over silhouette score optimisation was preferred. SHAP analysis revealed latent seasonal and appliance-specific patterns that were otherwise undetectable through conventional silhouette-based clustering methods. The integration of LLM with SHAP yielded interpretable energy consumption trends, which were validated through cluster centroid analysis, thereby demonstrating enhanced interpretability through the identification of factors not readily discernible in the raw data.

After analysing cluster behaviour using centroids, SHAP-based explainable AI identified the most influential factors affecting energy consumption across clusters, as shown in Figure 5. Cluster 0 exhibited balanced moderate positive impacts from mean temperature, cloud cover, global radiation, and max temperature, characterising moderate weather conditions with distributed feature influence. Cluster 1 demonstrated mean temperature dominance with the widest SHAP value spread, while other features clustered tightly around zero, indicating temperature-variable conditions where mean temperature serves as the primary distinguishing factor. Cluster 2 displayed the broadest SHAP value spreads for global radiation and mean temperature, with global radiation showing particularly high positive impacts, representing high solar radiation days. Cluster 3 showed high positive SHAP impacts from mean temperature and global radiation with tighter distributions than Cluster 2, indicating warm, sunny conditions of moderate intensity. Cluster 4 exhibited strong positive mean temperature

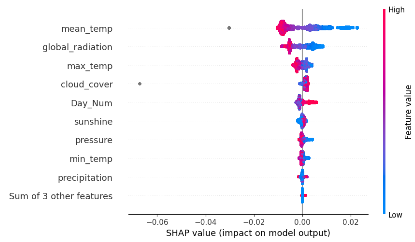

(a) Cluster 0 SHAP Values

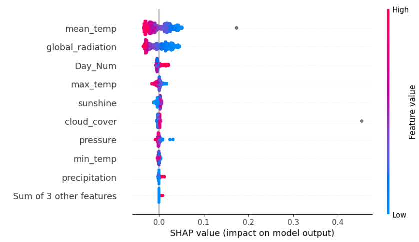

(b) Cluster 1 SHAP Values

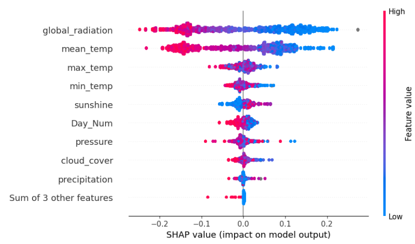

(c) Cluster 2 SHAP Values

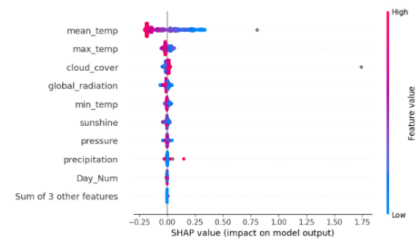

(d) Cluster 3 SHAP Values

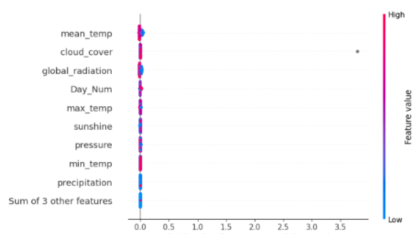

(e) Cluster 4 SHAP Values

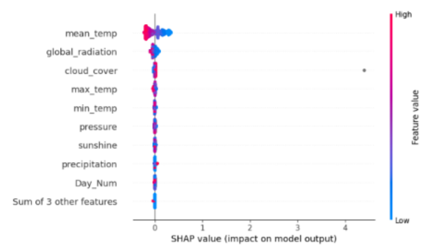

(f) Cluster 5 SHAP Values

Fig. 5: Impact of model outputs using SHAP values for clusters 0 to 5

impacts while other features clustered around zero with slight negative effects, demonstrating temperature-dominated warm conditions with minimal secondary weather influence. Cluster 5 displayed tight clustering around zero SHAP values across all features with only slight positive impacts from mean temperature and global radiation, representing neutral baseline weather conditions with no dominant factors. Table 1 presents a comparative summary of energy usage patterns across the identified clusters, generated through a hybrid approach that combines XAI and LLMs.

4.2 Discussion

This study addresses the critical limitations of prior methods by integrating deterministic K-means clustering for stable segmentation with SHAP-based ex-

Table 1: Energy Usage Cluster Characteristics and SHAP-Based Interpretations

Cl.	Peak (WD)	Peak (WE)	Lowest	Evening	Pattern	Type	SHAP Features	LLM-based SHAP Justification
0	6:00 PM	6:00 PM	3:00 AM	High at 6 PM	Rising to evening peak	Multi-Generational Households	**Temperature, Time of Use**	Consistent evening peaks with moderate weather sensitivity suggest large households with varied schedules but predictable dinner/evening routines
1	6:00 PM	6:00 PM	3:00 AM	Decline after 9 PM	Distinct AM/PM peaks	Remote Workers/Home-Based Professionals	**Mean Temperature**	Different weekend patterns indicate flexible schedules, while variable temperature response suggests active climate control during work-from-home periods
2	8–9 AM, 6 PM	3 PM, 9 PM	3–4 AM	High at 9 PM	Two clear peaks	Retirees/Empty Nesters	**Peak Periods, Heating/Cooling**	Morning and evening peaks with strong weather responsiveness indicate daytime occupancy and active use of natural lighting/heating
3	3–4 PM, 7 PM	3 PM, 6 PM	3–4 AM	Highest at 7 PM	Afternoon, family routine peaks	Traditional Families	**Time of Use, Seasonal Effects**	After-school and dinner peaks with balanced environmental response suggest family schedules with moderate energy consciousness
4	6:00 PM	6:00 PM	3:00 AM	Falls after 9 PM	Evening-dominant trend	Young Professionals /HVAC-Dependent Users	**Temperature, Weekday/Weekend**	Strong temperature dependence with work-schedule peaks suggests heavy reliance on heating/cooling systems, typical of apartment dwellers or poorly insulated homes
5	6 PM, 9 PM	8:00 PM	3–6 AM	High at 8 PM	Evening-focused	Vacation/Second Homes	**Peak Time, Seasonal Variation**	Low overall variation and weather independence suggest intermittent occupancy patterns typical of seasonal or secondary residences

plainability for transparent insights. Enhancing reliability even with sparse or noisy data. Our approach incorporates multidimensional clustering that combines environmental and temporal features, integrates daily, seasonal, and hourly analyses, and accounts for environmental influences and bank holidays, thereby producing more affluent consumer segments that reflect real-world dynamics and surpass traditional demographic-based segmentation. By leveraging GPT-4o powered contextualisation with domain-specific knowledge, our method mitigates biases from incomplete historical records or inaccurate demographic data. The hybrid methodology strikes a balance between quantitative accuracy and qualitative insights. Integrating explainable AI with contextualised reasoning enhances accessibility and interpretability for stakeholders of varying technical proficiency. This approach supports scalable recommendations and detailed consumer profiles, incorporating evidence-based behavioural reasoning that captures nuanced patterns often overlooked by purely quantitative analyses. The generated typical consumer types and their identified consumption patterns now enable informed decision-making and tailored reporting, supporting the development of targeted energy-saving strategies and customised tariffs for each consumer group. The framework advances existing methodologies by ensuring numerical robustness, enhanced interpretability, contextual adaptability, and improved accessibility, effectively bridging trade-offs between accuracy, explainability, and practical utility for energy profiling in smart buildings. This study was limited by the absence of detailed consumer behavioural data and energy tariff information in the knowledge base, which affects the quality of the generated recommendation report. Furthermore, reliance on advanced LLM models with larger context windows imposed computational and financial constraints.

5 Conclusion and Future Work

Understanding energy consumption data and identifying typical consumer patterns using explainable AI techniques integrated with LLMs marks a significant advancement in smart meter data analysis. The proposed hybrid methodology offers more profound insights into consumption behaviours than traditional methods alone. Profiling analysis revealed varying dependencies on environmental factors, particularly weather, providing a nuanced understanding of household energy use with broader implications for energy management. By contextualising technical findings through LLMs, complex consumption patterns become more accessible to non-expert stakeholders, enabling informed decision-making. The system's ability to generate personalised recommendations based on identified patterns and domain knowledge represents a significant step toward behaviour-oriented energy solutions. This solution could be integrated with IoT devices to respond to peak usage times for each category of consumers by adjusting HVAC systems, thereby enhancing operations in smart building systems. Demand response policies, in this context, can empower consumers by increasing awareness of how daily activities influence electricity usage. For energy providers and businesses, such policies support improved load modelling, better client understanding, and enhanced network stability. Additionally, insights into user behaviour may drive the emergence of innovative business models that monetise flexible energy. A comparative analysis with existing methods demonstrates the proposed approach's advantage in maintaining numerical stability while enhancing interpretability through contextual reasoning. Unlike purely data-driven or LLM-only models, this hybrid method strikes a balance between analytical precision and human-readable insights, thereby overcoming the limitations of each when applied independently. However, the recommendations currently generated may seem generalised due to limited depth in domain-specific content. Future research should aim to integrate a more comprehensive knowledge base with richer contextual data, such as occupant lifestyles, building characteristics, household size, and working patterns, to refine insights and evidence-based reasoning.

References

1. Agee, P., Gao, X., Paige, F., McCoy, A., Kleiner, B.: A human-centred approach to smart housing. Building research & information **49**(1), 84–99 (2021)
2. Bai, Y., Camal, S., Michiorri, A.: News and load: A quantitative exploration of natural language processing applications for forecasting day-ahead electricity system demand. IEEE Transactions on Power Systems **39**(5), 6222–6234 (2024)
3. Batalla-Bejerano, J., Trujillo-Baute, E., Villa-Arrieta, M.: Smart meters and consumer behaviour: Insights from the empirical literature. Energy Policy **144**, 111610 (2020)
4. Bobek, S., Kuk, M., Szelążek, M., Nalepa, G.J.: Enhancing cluster analysis with explainable ai and multidimensional cluster prototypes. IEEE Access **10**, 101556–101574 (2022)
5. DataStore, L.: Smartmeter energy consumption data in london households (2018), accessed: 2025-02-25

6. Du Toit, J., Davimes, R., Mohamed, A., Patel, K., Nye, J.: Customer segmentation using unsupervised learning on daily energy load profiles. J. Adv. Inf. Technol **7**(2), 69–75 (2016)
7. Gas, B.: Off-peak electricity times explained - is it cheaper at night? (2025), accessed: 2025-02-23
8. GOV.UK: Smart meters in great britain, quarterly update september 2024. https://www.gov.uk/government/statistics/smart-meters-in-great-britain-quarterly-update-september-2024 (2024), accessed: Feb. 25, 2025
9. Han, D., Bai, H., Wang, Y., Bu, F., Zhang, J.: Day-ahead aggregated load forecasting based on household smart meter data. Energy Reports **9**, 149–158 (2023)
10. Huang, Y., Luo, X.: A method based on lie group machine learning for multivariate time-series clustering. In: 2024 6th International Conference on Data-driven Optimization of Complex Systems (DOCS). pp. 638–643. IEEE (2024)
11. Ichmoukhamedov, T., Hinns, J., Martens, D.: How good is my story? towards quantitative metrics for evaluating llm-generated xai narratives. arXiv preprint arXiv:2412.10220 (2024)
12. Jain, R.K., Smith, K.M., Culligan, P.J., Taylor, J.E.: Forecasting energy consumption of multi-family residential buildings using support vector regression: Investigating the impact of temporal and spatial monitoring granularity on performance accuracy. Applied Energy **123**, 168–178 (2014)
13. Jin, M., Zhang, Y., Chen, W., Zhang, K., Liang, Y., Yang, B., Wang, J., Pan, S., Wen, Q.: Position: What can large language models tell us about time series analysis. In: Forty-first International Conference on Machine Learning (2024)
14. Joeng, H.C., Joo, S.K.: Personalized electricity tariff recommendation method incorporating behavioral changes using customer profiles in the absence of historical metering data. IEEE Access (2024)
15. Lavin, A., Klabjan, D.: Clustering time-series energy data from smart meters. Energy efficiency **8**, 681–689 (2015)
16. Lovmar, L., Ahlford, A., Jonsson, M., Syvänen, A.C.: Silhouette scores for assessment of snp genotype clusters. BMC genomics **6**(1), 35 (2005)
17. Marcílio, W.E., Eler, D.M.: From explanations to feature selection: assessing shap values as feature selection mechanism. In: 2020 33rd SIBGRAPI conference on Graphics, Patterns and Images (SIBGRAPI). pp. 340–347. Ieee (2020)
18. Osama, S., Alfonse, M., Salem, A.B.M.: An efficient algorithm for extracting appliance-time association using smart meter data. Heliyon **5**(8) (2019)
19. Project, E.: European climate assessment & dataset (eca&d) – daily data (2025), accessed: 2025-02-25
20. Tronchin, L., Cordelli, E., Celsi, L.R., Maccagnola, D., Natale, M., Soda, P., Sicilia, R.: Translating image xai to multivariate time series. IEEE Access (2024)
21. Van Zyl, C., Ye, X., Naidoo, R.: Harnessing explainable artificial intelligence for feature selection in time series energy forecasting: A comparative analysis of grad-cam and shap. Applied Energy **353**, 122079 (2024)
22. Wang, X., Feng, M., Qiu, J., Gu, J., Zhao, J.: From news to forecast: Integrating event analysis in llm-based time series forecasting with reflection. Advances in Neural Information Processing Systems **37**, 58118–58153 (2024)
23. Wastensteiner, J., Weiss, T.M., Haag, F., Hopf, K.: Explainable ai for tailored electricity consumption feedback–an experimental evaluation of visualizations. arXiv preprint arXiv:2208.11408 (2022)