

Article

An Empirically Validated Framework for Automated and Personalized Residential Energy-Management Integrating Large Language Models and the Internet of Energy

Vinícius Pereira Gonçalves ^{1,*}, Andre Luiz Marques Serrano ^{1,*}, Gabriel Arquelau Pimenta Rodrigues ^{1,*},
Matheus Noschang de Oliveira ¹, Rodolfo Ipolito Meneguette ², Guilherme Dantas Bispo ¹,
Maria Gabriela Mendonça Peixoto ¹ and Geraldo Pereira Rocha Filho ³

¹ Department of Electrical Engineering, University of Brasilia, Federal District, Brasilia 70910-900, Brazil; matheus.oliveira@redes.unb.br (M.N.d.O.); guilherme.bispo@redes.unb.br (G.D.B.); maria.peixoto@unb.br (M.G.M.P.)

² Institute of Mathematical and Computer Sciences, University of Sao Paulo, São Carlos 13566-590, Brazil; menegette@icmc.usp.br

³ Department of Exact and Technological Sciences, State University of Southwest Bahia, Vitória da Conquista 45083-900, Brazil; geraldo.rocha@uesb.edu.br

* Correspondence: vpgvinicius@unb.br (V.P.G.); andrelms@unb.br (A.L.M.S.); gabriel.arquelau@redes.unb.br (G.A.P.R.)

Abstract

The growing global demand for energy has resulted in a demand for innovative strategies for residential energy management. This study explores a novel framework—MELISSA (Modern Energy LLM-IoE Smart Solution for Automation)—that integrates Internet of Things (IoT) sensor networks with Large Language Models (LLMs) to optimize household energy consumption through intelligent automation and personalized interactions. The system combines real-time monitoring, machine learning algorithms for behavioral analysis, and natural language processing to deliver personalized, actionable recommendations through a conversational interface. A 12-month randomized controlled trial was conducted with 100 households, which were stratified across four socioeconomic quintiles in metropolitan areas. The experimental design included the continuous collection of IoT data. Baseline energy consumption was measured and compared with post-intervention usage to assess system impact. Statistical analyses included k-means clustering, multiple linear regression, and paired *t*-tests. The system achieved its intended goal, with a statistically significant reduction of 5.66% in energy consumption (95% CI: 5.21–6.11%, $p < 0.001$) relative to baseline, alongside high user satisfaction (mean = 7.81, SD = 1.24). Clustering analysis ($k = 4$, silhouette = 0.68) revealed four distinct energy-consumption profiles. Multiple regression analysis ($R^2 = 0.68$, $p < 0.001$) identified household size, ambient temperature, and frequency of user engagement as the principal determinants of consumption. This research advances the theoretical understanding of human–AI interaction in energy management and provides robust empirical evidence of the effectiveness of LLM-mediated behavioral interventions. The findings underscore the potential of conversational AI applications in smart homes and have practical implications for optimization of residential energy use.

Keywords: internet of things (IoT); large language models (LLMs); residential energy management; smart energy solutions; energy consumption optimization; machine learning for energy efficiency; user-centered energy management; smart home automation; energy consumption patterns; personalized energy recommendations



Academic Editor: Álvaro Gutiérrez

Received: 16 June 2025

Revised: 8 July 2025

Accepted: 9 July 2025

Published: 15 July 2025

Citation: Gonçalves, V.P.; Serrano, A.L.M.; Rodrigues, G.A.P.; de Oliveira, M.N.; Meneguette, R.I.; Bispo, G.D.; Peixoto, M.G.M.; Filho, G.P.R. An Empirically Validated Framework for Automated and Personalized Residential Energy-Management Integrating Large Language Models and the Internet of Energy. *Energies* **2025**, *18*, 3744. <https://doi.org/10.3390/en18143744>

Copyright: © 2025 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license

(<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Energy management represents a systematic and strategic approach to monitoring and controlling energy resources to improve efficiency and promote sustainability. This research encompasses several key practices, including the accurate measurement of energy consumption, the implementation of optimized strategies for electricity use, and the encouragement of the adoption of renewable energy sources. When applied effectively, energy management not only minimizes adverse environmental impacts but also leads to significant reductions in operational and household costs [1]. Growing global concern over climate change and the need to reduce carbon emissions have driven the development of innovative solutions for optimizing energy consumption across various sectors [2], with particular attention to the residential sector due to its considerable potential for savings.

Residential electricity consumption represents a substantial share of global energy demand, accounting for approximately 27% of demand in developed countries [3]. Urbanization and the proliferation of electronic devices have driven a 2.3% annual increase in residential consumption over the past decade [4]. However, many households face persistent inefficiencies due to suboptimal appliance usage, a lack of control mechanisms, and the absence of intelligent management tools [5]. Additionally, the complexity of energy data and the need for technical expertise further hinder the implementation of effective conservation strategies.

With the growing adoption of ubiquitous and pervasive computing, an unprecedented opportunity has emerged to fundamentally transform the way energy is used in residential contexts. Automated systems enable more efficient control of devices, whether through remote appliance management, automatic adjustments to consumption based on users' specific needs, or integration with renewable energy sources such as solar panels [6,7]. The connectivity provided by the IoT has been a key catalyst for this evolution, enabling devices to communicate with each other and respond to stimuli in real-time, thereby optimizing electricity use and significantly reducing waste [8]. Empirical studies demonstrate that households equipped with energy-management systems based on ubiquitous computing can achieve consumption reductions of between 15% and 30% [9,10], representing not only substantial financial savings but also a significant contribution to reducing carbon footprints [11].

Ubiquitous computing, characterized by the seamless integration of technology into everyday environments, offers unprecedented opportunities for energy monitoring and control. In these systems, sensors distributed throughout residential spaces continuously collect data on usage patterns, environmental conditions, and energy consumption, providing critical inputs for management systems [12]. Real-time data acquisition enables the identification of inefficient consumption patterns through time-series analysis [13], early detection of equipment failures via outlier-detection algorithms [14], automatic adaptation to environmental conditions and user preferences using reinforcement learning techniques [15], and continuous optimization through machine learning algorithms with incremental updates [16]. The integration of these capabilities within a unified framework marks a substantial evolution from traditional approaches to energy management, which typically rely on manual intervention and retrospective data analysis.

Within this context, the Internet of Energy (IoE) emerges as a transformative solution for enhancing energy efficiency and addressing issues related to uncontrolled consumption. IoE refers to the interconnection of intelligent energy devices within a digital ecosystem, facilitating comprehensive monitoring and the automation of processes [11,17,18]. The convergence of smart home technologies with advanced energy-management systems accelerates the modernization of conventional electrical grids, transforming them into smart grids that can distribute energy more efficiently and sustainably [19]. This integrated model

enables bidirectional monitoring of residential energy flows [20–22], automated responses to grid signals such as price fluctuations or availability [23], seamless integration of multiple energy sources [24], and predictive optimization based on sophisticated algorithms capable of anticipating consumption trends and environmental conditions [25]. Implementing IoE in residential settings fosters the emergence of “prosumers”—individuals who both consume and produce energy, thus redefining the relationship between households and the power grid [20,26].

Integrating the Internet of Energy (IoE) with LLMs marks a new paradigm in how households interact with and control their energy usage. LLMs offer the potential to transform the interpretation of energy data, making insights more accessible and actionable for end users [27]. These models can efficiently process data generated by smart devices, delivering personalized recommendations aligned with individual consumption patterns [28]. Instead of requiring users to interpret complex numerical reports, virtual assistants powered by large language models (LLMs) can provide detailed guidance on reducing waste and enhancing energy efficiency [27]. This paradigm shift offers a substantially more intuitive interface, empowering individuals without specialized technical expertise to easily understand and act upon energy-management insights [29]. Furthermore, LLMs enable natural and seamless communication between consumers and energy systems, supporting real-time responses to queries and personalized strategy recommendations [30].

Despite significant advances in residential energy-management systems and the growing application of LLMs across various domains, a substantial gap remains in the scientific literature regarding the effective integration of these technologies to optimize residential energy consumption [31–33]. However, empirical evidence underscores the effectiveness of LLM-based interfaces: a controlled study conducted by the University of Oxford found that households that used LLM-driven systems engaged in 27% more energy-conservation actions compared to those using traditional control-panel interfaces ($n = 150$, $p < 0.01$) [34]. Specifically, there is a lack of empirical studies that quantitatively assess the effectiveness of LLM-based interfaces for energy management in different socioeconomic contexts, identify key factors for successful technology implementation, analyze how different consumption profiles respond to recommendations generated by LLMs, and measure the impact of contextual variables on the effectiveness of these systems [33,35]. This gap is particularly relevant given the heterogeneity of residential energy consumption and the need for solutions that are adaptable to different socioeconomic realities [36,37].

This study aims to address this gap by proposing and empirically evaluating the MELISSA framework. Specifically, our objective is to address the following research questions: (1) What is the effectiveness of an LLM-based energy-management system in reducing residential energy consumption? (2) How does the system’s effectiveness vary across different socioeconomic classes and consumption profiles? (3) What are the main determinants of residential energy consumption, and how can an LLM-based system address them? (4) How do contextual variables (temperature, occupancy, special events) affect the system’s performance? By answering these questions, our aim is to contribute to the advancement of knowledge on intelligent energy-management systems and provide valuable insights for the development of more effective and inclusive solutions.

The remainder of this article is organized as follows. Section 2 presents a comprehensive literature review on home energy-management systems and the integration of IoT and LLMs. Section 3 details the methodology adopted for system implementation and evaluation. Section 4 discusses the empirical results, including consumption reduction, usage profiles, and determinants. Section 5 provides a critical discussion of the findings, and Section 6 concludes the study with final remarks and future research directions.

2. Literature Review

Home Energy-Management Systems (HEMS) represent a significant evolution in the interface between technology and domestic energy consumption. Historically, the development of these systems dates back to the 1990s, when the first energy-monitoring devices began to be implemented in residences, primarily in developed countries, as part of energy-efficiency initiatives [38]. The evolution of these systems has followed technological developments, progressing from simple smart meters to complex and interconnected ecosystems. According to Morales et al. (2022), the evolutionary trajectory of HEMS can be organized into three distinct generations: basic monitoring systems (1990–2005), automated control systems (2005–2015), and adaptive intelligent systems (2015–present) [39].

MELISSA belongs to the third generation of HEMS, which is characterized by intelligent, adaptive systems that incorporate machine learning and natural language interfaces. As shown in Figure 1, these features aim to simplify user interaction and broaden access to energy efficiency solutions. Given that residential buildings account for nearly 40% of global energy consumption, with the potential for a reduction of 20–30% through advanced management, the need for inclusive and scalable solutions has never been more critical [40].




Figure 1. Schematic representation of HEMS.

Based on the theoretical and technological foundations discussed above, this research works on the development of the MELISSA (Modern Energy LLM-IoE Smart Solution for Automation) framework, a Home Energy-Management (HEMS) solution that integrates the Internet of Things with Broad Language Models. Thus, MELISSA is designed to significantly enhance the interaction between individuals as source of generation energy systems, offering an evidence-based approach to optimizing energy consumption. The design of the framework was guided by principles of user-centered design, accessibility, and adaptability to different socioeconomic contexts, aiming to maximize its impact and adoption in diverse residential environments.

2.1. Integration Between IoT and LLMs: Theoretical Foundations

To operationalize the integration between IoT and LLM technologies, MELISSA employs a modular two-agent architecture. The first is the Data-Analyst Agent, responsible for low-level data acquisition, filtering, pattern recognition, and statistical modeling. The second is the Energy-Management Agent, which incorporates a fine-tuned large language model (LLM, GPT-3.5, 175B parameters) that serves as a natural language interface between the system and the user. Together, these components enable real-time analytics and personalized interaction through a conversational framework.

The integration between IoT and LLMs represents an emerging field with transformative potential for energy-management systems. This technological convergence is based

on the complementarity between the real-time data-collection capability provided by IoT devices and the advanced contextual processing offered by LLMs. As highlighted by Zhang et al. (2024), this integration allows for overcoming traditional limitations of IoT systems related to contextual interpretation and communication with users [41].

The IoT-LLM integration architecture, as implemented in MELISSA, is based on the federated edge computing paradigm, where data processing occurs close to the generation source, reducing latency and preserving privacy. This architectural model is particularly relevant for residential energy applications, where data sensitivity and real-time response needs are critical factors [42]. Recent studies demonstrate that implementing this architecture can reduce the system's energy consumption by up to 60% compared to cloud-centralized approaches, an aspect frequently overlooked in the development of HEMS solutions [43].

MELISSA employs a two-agent architecture that works together to process energy data and communicate insights to users, as illustrated in Figure 2. The first component, the Data-Analyst Agent, is responsible for processing raw data collected from IoT devices in the household. Its main functions include the collection and preprocessing of energy-consumption data using noise filtering and normalization techniques, the identification of consumption patterns and anomalies through outlier detection algorithms (DBSCAN, $\varepsilon = 0.5$, MinPts = 5), the condensation of large data volumes into concise summaries using dimensionality reduction techniques (PCA, explained variance > 95%), comparative analysis with historical and reference benchmarks through statistical tests (paired *t*-test, ANOVA), and the generation of energy performance metrics with 95% confidence intervals. This agent operates continuously in the background, processing real-time data streams and generating periodic analyses that are stored in a central database for later access.

This design ensures low-latency, privacy-preserving computation while enabling a seamless flow of contextual information to the LLM for user interaction. The modularity also supports component upgrades and behavioral feedback loops between analytics and the recommendation engine.




Figure 2. MELISSA architecture overview. The Data-Analyst Agent (left) processes real-time IoT data and feeds summaries into the Energy-Management Agent (right), which incorporates a large language model (LLM) to generate natural language recommendations and respond to user queries. Arrows represent data flow and bidirectional API communication between components.

2.2. Advanced Techniques for Energy-Data Processing

Efficient processing of energy data represents a significant challenge due to the noisy, multidimensional, and temporally dependent nature of these data. The techniques implemented in MELISSA reflect the state of the art in this field, combining traditional statistical approaches with machine learning methods. The application of the DBSCAN (Density-Based Spatial Clustering of Applications with Noise) algorithm with parameters $\epsilon = 0.5$ and $\text{MinPts} = 5$ for anomaly detection represents a methodological choice based on empirical evidence. Comparative studies demonstrate that this specific configuration offers an optimal balance between sensitivity and specificity in identifying anomalous patterns in residential energy data [44].

The use of dimensionality-reduction techniques, such as PCA (Principal Component Analysis) with an explained variance threshold greater than 95%, allows efficient condensation of large data volumes without significant information loss. This approach is particularly relevant in the energy IoT context, where data dimensionality can increase exponentially with the number of monitored devices. Recent research indicates that applying PCA to energy data can reduce storage volumes by up to 70% while maintaining the ability to reconstruct relevant patterns with greater than 95% precision [45].

The second component, the Energy-Management Agent, integrates a fine-tuned LLM (GPT-3.5) to interpret analytical results and interact with users. It performs real-time natural language generation with a response latency below 500 ms, contextualizes energy-usage patterns, and explains technical concepts in language adapted to the user's literacy level, using a classification model to determine the complexity of vocabulary and sentence structure. The LLM also tracks recommendation adherence and dynamically adjusts the choice of advised strategies.

The MELISSA framework achieves energy reduction by closing the loop between continuous data analysis and actionable, user-adapted feedback. The Data-Analyst Agent identifies inefficiencies in household consumption patterns, such as abnormal spikes during off-peak hours, overuse of appliances during warm periods, or inconsistent device scheduling. It does so using unsupervised learning (k-means for profiling, DBSCAN for anomalies) and statistical comparisons with historical data and reference benchmarks. These insights are interpreted by the Energy-Management Agent, which translates them into targeted, comprehensible recommendations using a fine-tuned LLM. Examples include suggesting adjustments to thermostat settings based on outside temperature, proposing automation routines aligned with low-tariff periods, and advising changes in appliance usage. The system prioritizes recommendations based on estimated cost-benefit ratios and user receptiveness, adapting future suggestions based on observed outcomes. This dynamic and personalized feedback mechanism enables MELISSA to drive sustained reductions in consumption while maintaining user engagement and autonomy.

2.3. MELISSA Engine: Modular Recommendation System

The MELISSA engine is responsible for generating personalized energy-efficiency recommendations by combining rule-based logic with data-driven insights. It is composed of five interoperable modules (M1 to M5), each designed to address a different layer of the recommendation pipeline, with these ranging from data ingestion to adaptive user messaging. This modular design ensures scalability, transparency, and the ability to accommodate heterogeneous user contexts.

Module M1: Consumption Profiling. This component continuously monitors consumption patterns via data collected from smart meters and plug-level sensors and occupancy data. It aggregates time-series data and performs clustering (e.g., k-means) to characterize

daily and weekly usage routines. The profiles serve as a baseline for identifying deviations or inefficiencies.

Module M2: Contextual Analysis. This module enriches energy-consumption data with contextual features such as outdoor temperature, solar irradiance, dynamic tariffs, and occupancy schedules. Through feature engineering and normalization, it constructs a multidimensional view of the user's operational environment, enabling the generation of more relevant recommendations.

Module M3: Inefficiency Detection. Using statistical comparisons and anomaly-detection algorithms (e.g., DBSCAN, time-series decomposition), this module identifies deviations from optimal behavior or energy-waste events. Examples include appliances left on during periods of absence or excessive heating during temperate weather.

Module M4: Recommendation Engine. Based on the outputs of M1–M3, this component selects suitable interventions from a predefined rulebase. Each rule is parameterized by device type, behavioral pattern, environmental condition, and expected impact. Recommendations are categorized into five macro-classes: behavioral tips, configuration changes, automation routines, equipment upgrades, and structural improvements. Each recommendation is scored based on its estimated energy-saving potential, complexity, and compatibility with user habits.

Module M5: Natural Language Personalization. This final module is responsible for translating technical recommendations into personalized, accessible messages. It integrates a fine-tuned large language model (GPT-3.5) that adapts the message tone and complexity based on user literacy and interaction history. Recommendations are enriched with contextual cues (e.g., "based on yesterday's afternoon usage") and explanatory components (e.g., "this could reduce your cooling cost by 15%") to improve clarity and engagement.

Together, these five modules operate as a dynamic recommendation system capable of producing actionable, user-centered suggestions tailored to each household's profile and context. The system is designed to function in a continuous learning loop, refining outputs as more data are gathered and user behavior evolves. The evaluation of this engine's effectiveness is presented in Section 4.6.

2.4. Operational Workflow of the Dual-Agent Architecture

The MELISSA framework relies on a two-agent architecture consisting of the Data-Analyst Agent and the Energy-Management Agent, which work in an integrated and asynchronous manner through a RESTful microservices architecture.

The Data-Analyst Agent is responsible for processing raw energy-consumption data collected from IoT devices installed throughout the household. It performs data cleaning, normalization, and synchronization of time-series inputs from multiple sensors. Using unsupervised algorithms such as DBSCAN, it detects anomalies and identifies unusual consumption events. Dimensionality-reduction techniques such as PCA are applied to summarize high-dimensional sensor data without significant information loss. The system also compares current usage patterns with historical baselines and statistical reference models using paired *t*-tests and ANOVA. Temporal usage patterns are extracted and clustered using k-means to generate consumption profiles, which are stored and forwarded to the Energy-Management Agent.

The Energy-Management Agent incorporates a fine-tuned GPT-3.5 language model and acts as the primary interface with the user. It interprets the results produced by the analyst agent and converts them into comprehensible natural language. The agent adjusts the complexity of explanations based on each user's digital literacy, as inferred by a textual-complexity classifier. Recommendations are then generated according to a cost-benefit strategy and contextual variables such as temperature, occupancy, and

tariff schedules. Users can query the system in natural language and receive instant feedback, with a response latency of less than 500 milliseconds. The agent also tracks the implementation of recommendations and may request new analyses from the analyst agent to refine its suggestions. This continuous feedback loop enables MELISSA to provide dynamic, personalized, and data-driven guidance to promote residential energy efficiency.

2.5. User-Centered Design and Socioeconomic Adaptability

The effectiveness of HEMS systems is intrinsically linked to their ability to adapt to the specific needs and contexts of users. MELISSA incorporates principles of user-centered design (UCD) that transcend mere usability, addressing aspects of cognitive accessibility and socioeconomic adaptability. As highlighted by Ribeiro and Providência (2021), the application of UCD in energy systems must consider not only functional aspects but also cultural, educational, and economic dimensions that influence users' relationship with technology [46].

Adaptability to different socioeconomic contexts represents a significant differentiating characteristic of MELISSA compared to conventional solutions. Studies indicate that the effectiveness of HEMS systems varies significantly across different socioeconomic strata, with adoption and engagement rates up to 60% lower in low-income populations when compared to high-income groups [47]. MELISSA addresses this disparity through adaptive mechanisms that adjust both interface complexity and the nature of energy recommendations according to the user's socioeconomic profile and level of technological literacy.

The communication between the two agents is bidirectional, allowing the management agent to request specific analyses from the analyst agent based on user interactions. This dual architecture enables the combination of rich information with the computational efficiency required for real-time processing. The system was implemented using a microservices architecture, with each agent operating as an independent service that communicates through RESTful APIs. This approach provides scalability and flexibility, enabling independent updates of components and the addition of new features without interrupting overall system operation.

To ensure compliance with data-protection regulations and safeguard user privacy, MELISSA was developed following privacy-by-design principles. All personally identifiable information (PII) was pseudonymized at the edge computing layer before transmission, in line with data protection regulations standards. Communication between IoT devices and the backend servers employed Transport Layer Security (TLS 1.3), and all data at rest were encrypted using the Advanced Encryption Standard (AES-256). The system design aligns with the General Data Protection Regulation (GDPR), incorporating principles such as data minimization, purpose limitation, and explicit user consent for data collection. Access to data was restricted using role-based access control (RBAC), with full audit trails and activity logging, following the guidelines of applicable regulations. These measures collectively ensured that user data remained secure, anonymized, and accessible only under strict operational controls.

2.6. Security and Privacy in IoT Energy Systems

The implementation of IoT systems in residential environments for energy monitoring raises significant concerns regarding data security and privacy. Energy-consumption patterns can reveal sensitive information about habits, routines, and even family composition, making the protection of these data a primary concern [48]. MELISSA implements a privacy-by-design approach, where security considerations are incorporated from the initial stages of development.

The two-agent architecture adopted by MELISSA significantly contributes to strengthening system security by creating a clear separation between the processing of raw data and the user interface. This functional segregation allows the implementation of granular access policies and anonymization mechanisms that limit the exposure of sensitive data [49]. Additionally, local processing and data-aggregation techniques are employed to minimize the transmission of potentially identifiable information to remote servers.

2.7. Environmental Impact and Sustainability

The contribution of HEMS systems such as MELISSA to sustainability objectives extends beyond the mere reduction of residential energy consumption. Studies indicate that effective implementation of these systems can result in carbon-emissions reductions of 5–15% per residence, depending on the local energy mix and consumption patterns [50,51]. This reduction is achieved not only through consumption optimization, but also by facilitating the integration of distributed renewable sources and participation in demand-response programs.

MELISSA incorporates life-cycle considerations in its design, minimizing the system's environmental impact through optimization of computational resources and careful component selection. This holistic approach to sustainability represents a significant advancement compared to conventional systems, which often overlook the environmental impact of their implementation and operation [52].

Therefore, the system uses a two-agent architecture (Data Analyst and Energy Management) that processes information from IoT devices and interacts with the user through natural language interfaces. The bidirectional arrows indicate real-time communication between the components, which allows for continuous adaptation and personalization based on user feedback.

3. Methodology

To evaluate the effectiveness of the MELISSA framework, we conducted a controlled experiment involving 100 households located in various urban and suburban areas and representing multiple socioeconomic levels. The study employed a 12-month longitudinal design, comprising an initial 3-month baseline phase during which energy-consumption data were passively collected without user interaction or activation of MELISSA's analytical or feedback features. This was followed by 9 months of active implementation of the MELISSA system, during which time all system functionalities were enabled. This approach enabled robust before-and-after comparisons, controlling for seasonal variations in energy consumption.

The selection of participating households was conducted through stratified sampling, ensuring adequate representation of diverse socioeconomic profiles, household sizes, and family compositions. Inclusion criteria were that households have a stable internet connection and at least five major electrical devices and that they agree to participate in the full 12-month study. Households that had already installed solar energy systems or that had planned significant structural changes during the study period were excluded. Strict data-privacy measures, including anonymization of personal data, end-to-end encryption for data transmission, and secure storage in ISO 27001-certified servers, were implemented.

3.1. Data Collection

Data collection was carried out through a combination of IoT sensors installed in the participating households and periodic questionnaires administered to the residents. The sensors included smart meters for monitoring total energy consumption, smart plugs for specific devices (such as refrigerators, air conditioners, washing machines, televisions, and

computers), temperature and humidity sensors in different rooms, and presence detectors to monitor occupancy patterns. Energy consumption data were continuously collected at 5-minute intervals, resulting in approximately 105,120 data points per household from the study period. Additionally, contextual data such as external temperature, humidity, and weather events were obtained from local weather stations. Information on energy tariffs was updated monthly from local utilities.

3.2. System Implementation

The implementation of MELISSA in the participating households followed a standardized protocol with five stages: (1) installation of physical infrastructure (sensors, communication hubs, displays); (2) configuration of software and integration with existing devices; (3) a two-week calibration period to establish initial consumption profiles; (4) full activation of the system with all functionalities; and (5) continuous monitoring and remote updates when necessary.

The system was configured to provide three main types of interactions with users: (a) automated daily reports summarizing energy consumption and highlighting opportunities for savings; (b) real-time alerts signaling anomalies or immediate optimization opportunities; and (c) a conversational interface for specific queries initiated by the user. Participants received initial training on how to interact with the system but were encouraged to use it according to their individual preferences.

During the implementation period, software updates were carried out remotely to fix bugs, improve analysis algorithms, and expand the LLM's knowledge base with newly identified patterns. All interactions between users and the system were recorded for later analysis, with the explicit consent of the participants.

3.3. Data Analysis

The analysis of the collected data was conducted using a mixed approach that combined traditional statistical methods with advanced machine learning techniques. To evaluate the overall effectiveness of MELISSA in reducing energy consumption, we compared the average daily consumption during the pre- and post-implementation periods using paired *t*-tests, controlling for seasonal variations through degree-day normalization.

To identify different energy-consumption profiles, we applied cluster analysis using the k-means algorithm to consumption data disaggregated by device and time of day. The optimal number of clusters was determined using the elbow method and validated with the silhouette coefficient. Each cluster was characterized in terms of temporal usage patterns, dominant devices, and associated demographic characteristics.

The determinants of energy consumption were investigated through multiple regression models, with daily consumption as the dependent variable and factors such as household size, number of occupants, outdoor temperature, day of the week, and presence of specific devices as independent variables. Multicollinearity was assessed using the variance inflation factor (VIF), and appropriate transformations were applied when necessary.

The effectiveness of the recommendations generated by MELISSA was evaluated through time-series analysis, comparing energy consumption before and after each recommendation was implemented. ARIMA models were used to control for temporal trends and seasonality, isolating the specific effect of each intervention. Additionally, sentiment analysis was applied to the textual interactions between users and the system to assess receptiveness to the recommendations.

Finally, we conducted a subgroup analysis to investigate how the effectiveness of MELISSA varied across different socioeconomic profiles and types of residences. Mixed-

effects models were used to accommodate the hierarchical structure of the data and identify significant interactions between user characteristics and energy-savings outcomes.

All statistical and machine learning analyses described in this section were implemented in Python 3.10 using the sci-kit-learn, stats models, NumPy, and pandas libraries. The models were executed on a workstation running Ubuntu 22.04 LTS that was equipped with an Intel Core i7-11700 processor (2.5 GHz, 8 cores) and 32 GB of RAM. This computational setup ensured efficient processing of the approximately 10 million data points collected during the study and is sufficient to reproduce the analysis pipeline described herein.

To ensure robustness and reproducibility of the machine learning models, we adopted the following training and validation procedures:

- DBSCAN was tuned via grid search with $\epsilon \in [0.3, 0.7]$ and $\text{MinPts} \in [4, 10]$, selecting the configuration with the highest silhouette coefficient.
- k-means clustering was validated using both the elbow method and the silhouette coefficient ($k = 4$, silhouette = 0.68). Clusters were characterized based on temporal device usage and demographic attributes.
- Multiple linear regression included multicollinearity checks using the Variance Inflation Factor (VIF). Variables with $\text{VIF} > 5$ were iteratively removed, and residual diagnostics were performed to validate assumptions of linearity.
- ARIMA models used for intervention analysis were selected via the `auto_arima` function from the `pmdarima` library, optimizing the Akaike Information Criterion (AIC) under seasonal constraints. Models were validated with time-series train-test splits.

All models were implemented using Python 3.10 with the `scikit-learn`, `statsmodels`, and `pmdarima` libraries.

4. Results

4.1. Methodological Justification for Household Exclusions

The scientific integrity of this study relies on the quality and consistency of the data analyzed. Although 100 households were initially enrolled in the MELISSA study, only 97 were retained for the final analysis. The exclusion of three units was based on established methodological principles:

Data inconsistencies and technical issues. Data collected via IoT sensors are susceptible to transmission errors, device failures, or inconsistent records. In two households, persistent technical problems resulted in missing or unreliable data across several key variables, rendering them unsuitable for longitudinal analysis.

Participant withdrawal. One household chose to withdraw from the study within the first month for personal reasons unrelated to the MELISSA framework.

Compliance with inclusion criteria. In longitudinal studies, maintaining homogeneity in the sample is essential. Households that undergo significant structural or behavioral changes—such as renovations, new energy systems, or occupancy changes—may introduce confounding variables. Although this did not apply directly to the three excluded homes, it remains a relevant criterion for exclusion in similar contexts.

In sum, the final sample of 97 households reflects those with complete, consistent, and reliable data over the full 12-month observation window. This approach ensures the validity and robustness of the study's findings.

4.2. Data Analysis

The analyzed sample comprised 97 households that remained actively in the study for the full 12 months, resulting in a high retention rate of 97%. The households exhibited considerable diversity in both physical and demographic terms, underscoring the heterogeneity of the population being studied. It was observed that the size of the residences

varied considerably, ranging from compact apartments of 45 m² to larger houses of 230 m², with an overall average of 112.3 m² and a standard deviation of 42.7 m², reflecting different housing profiles.

The average number of occupants per household was 3.2 individuals (SD = 1.4), with configurations ranging from single-person households to those with five or more occupants. This demographic diversity, detailed in Table 1, indicates that the study included a representative sample of different family arrangements.

Regarding socioeconomic status, the households were relatively evenly distributed among income quartiles, with a slight predominance of households in the intermediate ranges. The respondent profiles revealed an average age of 42.7 years (SD = 13.5 years), with a relatively balanced gender composition: 55% women and 45% men. Regarding educational level, there was a predominance of individuals with higher education (38% with a bachelor's degree and 15% with a postgraduate degree), although participants with elementary and high-school education were also included, ensuring educational diversity.

In terms of energy infrastructure, the collected data indicate a broad availability of equipment that consumes a significant amount of energy: most households had air conditioning (87%), refrigerators (100%), and televisions (100%), and 93% had washing machines. The average number of units per household also demonstrates the prevalence of having multiple devices in the same home, as seen with televisions, resulting in a significant increase per household. Finally, the average energy consumption during the baseline period was 312.7 kWh/month (SD = 156.3 kWh/month), showing high variability across households, with values ranging from 98.4 kWh/month to 723.5 kWh/month. This range reflects different consumption habits, installed infrastructure, and socioeconomic conditions—key aspects that were analyzed throughout the study in relation to energy-management strategies.

Table 1. Characteristics of the participating households.

Item	Description
Number of households	97 (retention rate: 97%)
Average household size	112.3 m ² (SD = 42.7 m ²)
Size variation	From 45 m ² to 230 m ²
Average number of occupants	3.2 (SD = 1.4)
Occupant distribution	18% with 1 occupant; 35% with 2 occupants; 27% with 3–4 occupants; 20% with 5 or more occupants
Socioeconomic distribution	22% first quartile; 31% second quartile; 28% third quartile; 19% fourth quartile
Average respondent age	42.7 years (SD = 13.5 years)
Gender of respondents	55% women, 45% men
Educational level of respondents	15% elementary education; 32% high school; 38% bachelor's degree; 15% postgraduate
Energy infrastructure	87% with air conditioning (1.8 units/household); 100% with refrigerator (1.1 units/household); 93% with washing machine; 100% with television (2.3 units/household)
Average energy consumption	312.7 kWh/month (SD = 156.3 kWh/month)
Range of energy consumption	From 98.4 kWh/month to 723.5 kWh/month

4.3. Impact on Energy Consumption Reduction

The implementation of the MELISSA framework led to a meaningful and consistent reduction in residential energy consumption, demonstrating not only the technical feasibility of intelligent energy management but also its practical effectiveness in real-world settings. After controlling for seasonal fluctuations and relevant contextual vari-

ables—including climate, household occupancy, and baseline consumption patterns—we observed an average monthly reduction of 5.66% in energy use (95% CI: 5.21–6.11%) during the intervention period when compared to the pre-intervention baseline. This decrease was statistically significant ($t(96) = 14.72, p < 0.0001$), corresponding to a mean savings of 17.7 kWh per household per month. This result suggests there is considerable potential for aggregate energy efficiency and carbon mitigation if this intervention were to be scaled to a broader population.

According to the Table 2, it is evident that “MELISSA” not only modifies consumption behavior in the short term, but also promotes long-term awareness and the continuous adoption of energy-saving practices. These findings underscore the importance of integrating technological interventions with behavioral approaches to achieve lasting environmental and economic benefits in the residential context.

Table 2. Reduction in energy consumption after implementation of the MELISSA framework.

Period	Reduction (%)	Acum (%)	95% CI	(kWh/Month)
Overall (12 months)	5.66%		5.21–6.11%	17.7
First Quarter	3.21%		2.87–3.55%	–
Second Quarter	5.94%		5.43–6.45%	–
Third Quarter	–1.32%	7.83%	7.21–8.45%	–

More importantly, the effect of MELISSA proved to be cumulative and adaptive over time, indicating a process of behavioral assimilation and growing user engagement. In the first quarter post-implementation, the average reduction was 3.21% (95% CI: 2.87–3.55%), already a meaningful initial impact. This improvement intensified in the second quarter, reaching 5.94% (95% CI: 5.43–6.45%), reflecting both the refinement of system recommendations and increased user familiarity. By the third quarter, the cumulative reduction reached 7.83% (95% CI: 7.21–8.45%), representing a progressive increase of 4.62 percentage points from the initial phase.

Figure 3 illustrates the trajectory of average energy consumption over the study period, highlighting the downward trend after the MELISSA implementation. Notably, the magnitude of the reduction increased over time, suggesting a cumulative effect as users became more familiar with the system and implemented more recommendations.

The shaded area represents the baseline period, while the dashed vertical line indicates the start of MELISSA implementation. The annotations show the average percentage reductions observed in each quarter following implementation.

The device-wise analysis revealed that the largest percentage reductions were observed in the consumption of energy for air conditioning (9.72%, 95% CI: 8.95–10.49%) and lighting (8.34%, 95% CI: 7.76–8.92%). Moderate reductions were observed for televisions and entertainment equipment (5.21%, 95% CI: 4.78–5.64%) and computers (4.87%, 95% CI: 4.32–5.42%). The reductions were smaller for refrigerators (2.13%, 95% CI: 1.87–2.39%) and washing machines (2.76%, 95% CI: 2.34–3.18%), reflecting the lower optimization potential for these essential devices.

The long-term effectiveness of MELISSA is rooted in its ability to foster awareness and promote sustained behavioral change. This is achieved through a built-in adaptive feedback mechanism that continuously learns from user interactions and consumption patterns. As users implement recommendations, the system tracks their responses and adjusts future suggestions accordingly—modifying tone, specificity, and timing to align with individual preferences and behaviors. Moreover, the LLM component contextualizes energy usage through personalized narratives, transforming technical insights into understandable, relatable stories. These narratives not only explain energy-consumption trends, but also

quantify the benefits of actions taken, reinforcing user learning and building a sense of ownership over energy decisions. By combining real-time feedback with behavioral reinforcement, MELISSA transitions users from one-time adjustments to a sustained pattern of proactive energy management.




Figure 3. Trajectory of Average Energy Consumption Throughout the Study.

4.4. Energy-Consumption Profiles

The cluster analysis identified four distinct energy-consumption profiles among the participating households, with a silhouette coefficient of 0.68 indicating good separation between clusters. Figure 4 presents the main characteristics of each identified profile.




Figure 4. Energy-consumption profiles identified by cluster analysis.

The graph shows the daily consumption patterns for the four identified clusters, with shaded areas representing different times of day. The annotations highlight the distinctive characteristics of each profile.

Cluster 1 ($n = 23$, 23.7% of the sample), named “Intensive Daytime Consumption”, was characterized by high consumption during the day (9 AM–5 PM) and intensive use of air conditioning and office equipment. This cluster predominantly included households

where at least one resident worked from home (78.3% of households in this cluster). The average monthly consumption for this group was 378.4 kWh (SD = 67.2 kWh).

Cluster 2 (n = 31, 32.0%), named “Moderate Nighttime Consumption”, had consumption peaks in the evening (6 PM–11 PM), with moderate use of television, lighting, and small appliances. This cluster mainly included families with school-aged children (74.2%) and both parents working outside the home (83.9%). The average monthly consumption was 287.6 kWh (SD = 52.8 kWh).

Cluster 3 (n = 19, 19.6%), named “High Constant Consumption”, was characterized by high and relatively constant consumption throughout the day, with intensive use of multiple devices simultaneously. This cluster comprised larger homes (average size of 156.7 m²) with more occupants (average of 4.7 people) and higher household incomes. The average monthly consumption was 512.3 kWh (SD = 98.4 kWh).

Cluster 4 (n = 24, 24.7%), named “Efficient Low Consumption”, showed consistently low consumption throughout the day, with limited use of high-consumption devices. This cluster primarily consisted of smaller homes (average size of 78.3 m²) with fewer occupants (average of 1.8 people) and behaviors already oriented towards energy efficiency. The average monthly consumption was 174.5 kWh (SD = 43.6 kWh).

Table 3 provides a consolidated summary of the main attributes of the four energy-consumption clusters identified through the analysis. It includes the sample size, relative proportion, distinguishing behavioral and household characteristics, and average and standard deviation of monthly energy consumption for each group. This tabular representation complements the visual patterns presented in Figure 4, allowing for a more detailed comparison between clusters and reinforcing the heterogeneity in energy usage across different household profiles.

Table 3. Summary of energy-consumption clusters.

Cluster	n	% of Sample	Characteristics	Average Consumption (kWh/Month)	Standard Deviation (kWh)
Cluster 1	23	23.7%	High daytime consumption (9 AM–5 PM), AC/office, remote work (78.3%)	378.4	67.2
Cluster 2	31	32.0%	Nighttime peaks (6 PM–11 PM), TV/lights/small appliances, school-aged children (74.2%), parents out (83.9%)	287.6	52.8
Cluster 3	19	19.6%	High and constant consumption, multiple appliances, larger homes (156.7 m ²), more occupants (4.7)	512.3	98.4
Cluster 4	24	24.7%	Low and constant consumption, few high-consumption appliances, smaller homes (78.3 m ²), fewer occupants (1.8)	174.5	43.6

The effectiveness of MELISSA varied significantly between the different clusters. The most significant percentage reductions were observed in Cluster 1 (7.82%, 95% CI: 7.13–8.51%) and Cluster 3 (6.94%, 95% CI: 6.27–7.61%). Cluster 2 showed a moderate reduction (5.23%, 95% CI: 4.78–5.68%), while Cluster 4 showed the smallest reduction (2.87%, 95% CI: 2.41–3.33%), likely due to its already low consumption level and pre-existing efficient practices.

4.5. Determinants of Energy Consumption

The estimated coefficients for the main predictors are presented in Figure 5, organized by absolute magnitude to facilitate intuitive comparison.

The multiple regression model identified a range of statistically significant determinants of residential energy consumption, achieving strong explanatory performance, with an adjusted coefficient of determination of $R^2 = 0.68$ and a root mean square error (RMSE) of 24.3 kWh. These metrics indicate that the model is not only statistically robust but also practically relevant for understanding the complex interplay between household characteristics and energy use.

In the graph, predictors are ordered from the most to the least impactful in terms of effect size. Positive coefficients indicate that the variable is associated with increased energy consumption, while negative coefficients suggest a tendency toward reduced energy consumption. The statistical significance of each predictor is confirmed by the associated p -values, which represent confidence in the model's findings and its ability to discern meaningful patterns.




Figure 5. Determinants of residential energy consumption.

Among all variables considered, house size emerged as the most influential factor: each additional square meter of floor space contributed, on average, 0.45 kWh/month ($p < 0.0001$). This association likely reflects the greater number of rooms, an increased surface area requiring climate control, and the higher probability of housing energy-intensive appliances. Similarly, the number of occupants showed a strong positive effect, with each additional resident accounting for 27.8 kWh/month ($p < 0.0001$), likely due to compounded energy use for lighting, electronics, and water heating.

Ambient temperature also had a notable impact: for every degree Celsius above the defined comfort temperature (22 °C), there was an estimated increase of 7.64 kWh/month in consumption ($p < 0.0001$). This finding underscores the sensitivity of household energy demand to climatic conditions, a relationship primarily driven by the intensive use of air-conditioning systems in warmer periods.

Importantly, socioeconomic factors added a deeper layer of nuance to the analysis. Higher household income was positively associated with energy consumption ($\beta = 0.23$, $p < 0.001$), reflecting patterns of more prevalent ownership of electronic devices and

comfort-related appliances. In contrast, a higher educational level was inversely related to consumption ($\beta = -0.18$, $p < 0.01$), suggesting that individuals with greater educational attainment may adopt more energy-efficient behaviors, independent of their income level.

Taken together, these findings reveal that residential energy consumption is driven not by a single factor, but rather by a complex interaction between physical infrastructure, household composition, environmental conditions, and socioeconomic context. Understanding these interdependencies is essential for designing more effective energy policies and interventions that are both socially equitable and environmentally sustainable.

Interestingly, the analysis revealed significant interactions between variables, as illustrated in Figure 6. Certain appliances significantly increase household electricity use: each air-conditioning unit adds, on average, 83.7 kWh/month ($p < 0.0001$); a clothes dryer, 42.3 kWh/month ($p < 0.0001$); an electric heater, 38.6 kWh/month ($p < 0.0001$); and a pool equipped with a pump, 35.2 kWh/month ($p < 0.001$). These figures reflect the energy-intensive nature of climate-control devices and continuous thermal or hydraulic processes.




Figure 6. Interactions among determinants of energy consumption.

However, the building and environmental context modulate this impact. When the dwelling offers good thermal insulation—double walls, ventilated roof, low-emissivity glazing—the effect of outdoor temperature on consumption drops sharply (interaction $p < 0.001$). Shading trees provide a complementary benefit: besides lowering the direct thermal load, they improve the micro-climate around the building, smoothing demand peaks ($p < 0.01$).

From a sociodemographic perspective, the total number of occupants is positively correlated with expenditure; however, this relationship becomes stronger when teenagers are present ($p < 0.01$). This age group typically uses electronic devices intensively—including

video games, streaming, and high-performance computers—which prolongs device run-time and increases indoor heat, triggering additional air-conditioner use.

The heatmap synthesizes these dynamics: cells in warmer tones signal strong positive correlations, whereas cooler tones indicate negative or protective effects. Outlined cells mark statistically significant interactions identified by the model, highlighting that energy efficiency arises from the interplay among technological choices, building characteristics, and everyday behaviors. This panorama highlights the importance of integrated policies that combine incentives for structural improvements, education for conscious consumption, and clear appliance labeling so that potential gains are translated into tangible savings for both users and the power grid.

4.6. Efficacy of Personalized Recommendations

Over the 18-month observation window, the MELISSA engine analyzed high-resolution consumption data and produced 12,483 personalized recommendations for the 97 participating households, averaging 128.7 suggestions per dwelling. Leveraging a hybrid rule-based/machine learning pipeline, these insights were clustered into five macro-classes, with pedagogical labels used to facilitate user comprehension:

- Immediate behavioral adjustments (42.3%): low-effort actions such as switching off standby devices or moderating thermostat set-points.
- Device-configuration optimizations (27.8%): fine-tuning schedules, power modes, and sensor thresholds in existing appliances.
- Equipment-replacement suggestions (12.4%): cost–benefit analyses advocating substitution by high-efficiency models certified A+++.
- Structural retrofits (8.7%): deeper measures, e.g., insulating the attic or installing double-glazed windows, to address envelope losses.
- Scheduled automations (8.8%): IFTTT-style routines that synchronise operation with tariff valleys or photovoltaic generation peaks.

Implementation rates diverged markedly across classes, as depicted in Figure 7. Behavioral tips were most readily adopted (73.2%), reflecting their negligible cost and immediate feedback. Configuration changes followed (68.7%), aided by MELISSA's step-by-step in-app wizards. Scheduled automation achieved 61.4% uptake, with this change limited mainly by users' familiarity with smart-home hubs. Capital-intensive recommendations saw lower adherence: only 32.1% of proposed equipment upgrades and a mere 17.3% of structural retrofits were executed, underscoring classical barriers of upfront investment, perceived payback, and contractor availability. These patterns echo the Technology Acceptance Model, where perceived effort and cost strongly mediate the translation of informational nudges into concrete energy-saving actions.

Time-series analysis revealed that implemented recommendations resulted in measurable reductions in energy consumption. On average, each implemented behavioral recommendation resulted in a reduction of 0.87 kWh/month (95% CI: 0.76–0.98 kWh/month); configuration optimizations resulted in a reduction of 1.34 kWh/month (95% CI: 1.18–1.50 kWh/month); scheduled automation resulted in a reduction of 1.76 kWh/month (95% CI: 1.54–1.98 kWh/month); equipment replacements resulted in a reduction of 4.23 kWh/month (95% CI: 3.87–4.59 kWh/month); and structural modifications resulted in a reduction of 6.78 kWh/month (95% CI: 6.12–7.44 kWh/month).

Sentiment analysis of interactions between users and the system revealed that recommendations presented with contextualized explanations and quantified benefits had a significantly higher probability of being implemented (OR = 2.34, 95% CI: 2.08–2.63, $p < 0.0001$). Additionally, recommendations referring to specific patterns observed in the

household were better received than generic recommendations (OR = 1.87, 95% CI: 1.65–2.12, $p < 0.0001$).




Figure 7. Efficacy of personalized recommendations. (A) Implementation rates by recommendation category show higher adherence to behavioral changes and lower rates of implementation for structural modifications. (B) The average reduction in energy consumption per implemented recommendation, with error bars representing 95% confidence intervals.

4.7. User Satisfaction and Experience

Before presenting the findings on user satisfaction, it is essential to highlight the ethical considerations associated with conducting research involving human participants. To address these concerns, the following documents were prepared and signed by all participants: (1) a Free and Informed Consent Form and (2) an Authorization Form for the Capture and Use of Images, Sounds, and Names.

The Free and Informed Consent Form clearly explained the study's objectives, its scientific nature, and the voluntary aspect of participation. The Authorization Form detailed the possibility of recording images and voices and assured participants that all collected data would be used exclusively for scientific purposes.

Figure 8 shows the progressive increase in user satisfaction during the nine months of MELISSA's implementation. The shaded area represents the 95% confidence interval, and the dashed line indicates the overall mean of 7.81/10. The boxes highlight the highest- and lowest-rated aspects of the system.

The aspects of the system most positively rated by users were the ease of understanding the information (8.43/10, SD = 0.97), the relevance of personalized recommendations (8.12/10, SD = 1.18), and the ability to respond to specific questions (7.96/10, SD = 1.32). The aspects with relatively lower ratings included the accuracy of savings predictions (7.21/10, SD = 1.56) and integration with existing devices (7.34/10, SD = 1.48).

The qualitative analysis of semi-structured interviews revealed several recurring themes. Participants frequently mentioned the “naturalness of interaction” as a positive feature of the system, with comments like: “It’s like talking to someone who really understands my home” and “I don’t need to decipher complicated graphs, the system explains everything in simple language.” Many participants also highlighted the “continuous learning” provided by the system: “I learned a lot about how my house consumes energy and now I make more conscious decisions.”

Reported challenges included occasional “understanding failures” when queries were too complex or ambiguous and “repetitive recommendations” in some cases. Some participants in smaller households reported that “the potential savings seemed limited” after the initial interventions with the greatest impact.




Figure 8. Evolution of user satisfaction over time.

5. Discussion

This study provides statistically supported evidence of MELISSA's effectiveness in residential energy optimization. The observed 5.66% average reduction, while moderate in relative terms, would translate into meaningful absolute savings if it were applied at scale. For example, national-level implementation in Brazil could save an estimated 8.7 TWh annually—equivalent to the generation capacity of a mid-sized hydroelectric facility.

The variation in system effectiveness across different consumption profiles offers valuable insights for future implementations. The fact that households with “Intensive Daytime Consumption” and “High Constant Consumption” patterns showed the largest percentage reductions suggests that MELISSA is particularly effective for users with greater optimization potential. This contrasts with traditional energy-efficiency approaches, which often achieve higher percentage savings in households with low energy consumption. This characteristic of MELISSA can be attributed to its ability to identify and address specific inefficiencies in complex consumption patterns, which are more prevalent in households with high consumption levels.

The analysis of the determinants of energy consumption confirms the importance of structural factors such as house size and number of occupants but also highlights the significant role of behavioral and contextual variables. The strong influence of ambient temperature, moderated by structural characteristics such as thermal insulation, emphasizes the need for adaptive solutions that dynamically respond to environmental conditions. MELISSA demonstrated the ability to provide contextualized recommendations that account for these complex interactions, a significant advancement over systems that treat variables in isolation.

The variation in implementation rates across different recommendation categories reflects the well-known practical and economic barriers commonly cited in the energy-conservation literature. Behavioral recommendations, which typically do not require financial investment, had the highest adoption rates, while structural modifications, which often involve significant costs and inconvenience, had more limited implementation. However, it

is noteworthy that even high-cost recommendations achieved implementation rates higher than those reported in previous studies with conventional interfaces. This suggests that the contextualized and personalized presentation of information through large language models (LLMs) can increase users' willingness to invest in energy-efficiency improvements.

The positive evolution of user satisfaction over time contrasts with the "feedback fatigue" pattern often observed in conventional energy-monitoring systems, where engagement typically decreases after an initial period of enthusiasm. MELISSA's ability to maintain and even increase engagement can be attributed to several factors: the conversational interface, which reduces the cognitive load associated with interpreting technical data; the continuous personalization, which increases the relevance of interactions; and the ability to respond to specific queries, which provides continuous value to the user.

The qualitative results from the interviews corroborate this interpretation, with participants frequently mentioning the "naturalness of interaction" as a positive feature. This finding aligns with previous human–computer–interaction research, which shows that interfaces that minimize the conceptual distance between the user's mental model and the system's representation result in greater satisfaction and effectiveness. MELISSA's LLM-based approach effectively reduces this conceptual distance by translating complex technical data into comprehensible narratives contextualized to the user's specific reality.

Although MELISSA's architecture integrates multiple components—IoT-based data acquisition, statistical and machine learning analysis, and an LLM-powered interface—our findings suggest that the LLM interface was the most decisive element for achieving effective reductions in energy consumption. This is primarily due to its capacity to translate complex analytical outputs into personalized, actionable, and easy-to-understand recommendations, which significantly boosted user adherence. While the IoT and data-analysis layers ensured robust pattern detection and system accuracy, it was the natural language interface that bridged the gap between technical insights and user behavior, as reflected in the high rates of engagement and recommendation implementation. Future modular studies should quantify these individual contributions more precisely, but preliminary evidence underscores the role of the LLM interface as a critical driver of impact.

6. Conclusions

This study provides compelling empirical evidence for the effectiveness and transformative potential of MELISSA, a novel framework that merges the capabilities of Internet of Things (IoT) technologies with Large Language Models (LLMs) to support sustainable residential energy management. Through personalized, natural language-based interactions, MELISSA demonstrated the ability to foster meaningful energy savings, reducing household consumption by an average of 5.66% over a 12-month period. Beyond numerical gains, this outcome represents a significant advancement in aligning everyday household behavior with broader goals of energy efficiency and environmental responsibility.

The significance of these results lies not only in the immediate reductions observed but also in the evidence of MELISSA's ability to encourage long-term shifts in consumption habits. The cumulative improvement over time suggests that the system acts as both a feedback mechanism and a learning facilitator, empowering users with contextualized and intelligible recommendations that support sustainable decision-making. Moreover, the system's performance across diverse socioeconomic profiles indicates a degree of scalability and inclusivity rarely achieved by traditional energy-management systems.

Nevertheless, this work acknowledges its limitations. The study was conducted in a single geographic region, which may limit its external validity in climates or cultural contexts that differ substantially. Additionally, while the 12-month period provided valuable insights into seasonal dynamics and medium-term adaptation, it remains insufficient to

assess the long-term resilience of behavioral changes or the system's response to disruptive household events, such as renovations or demographic shifts.

Furthermore, due to MELISSA's integrated architecture—encompassing sensors, data analytics, and an LLM-powered conversational interface—it was not possible to disentangle the contribution of each component to the overall impact. Future studies employing modular or factorial experimental designs will be crucial in identifying which components are most effective and under what conditions they are effective.

This study also did not fully explore the psychological dimensions that may underpin behavioral change, such as perceived efficacy, environmental concerns, or social norms. Integrating validated psychometric instruments into future research could significantly enhance the explanatory power of the findings and inform the refinement of user-engagement strategies.

One of MELISSA's strengths lies in its potential for generalization and adaptability to diverse international contexts. The system's modular architecture enables flexible integration with various smart devices, energy tariffs, and household infrastructures. To accommodate cultural and linguistic diversity, the LLM can be fine-tuned in multiple languages and aligned with regional communication norms. From a climatic perspective, MELISSA can be calibrated with local weather patterns, comfort-temperature ranges, and seasonal behaviors. Additionally, the system supports regulatory customization, allowing compliance with distinct privacy, data protection, and energy-efficiency policies in different jurisdictions. These characteristics position MELISSA as a scalable solution for global residential energy management, pending contextual customization efforts that are both feasible and technically supported by its current design.

Looking ahead, promising directions include integrating MELISSA with residential renewable-energy systems, enabling the optimization of self-consumption and intelligent grid interactions. Extending the framework to include other resources, such as water and gas, could transform MELISSA into a comprehensive platform for multisectoral household resource management. Moreover, advances in predictive modeling may allow the system to proactively detect and respond to anomalous consumption patterns, enhancing its capacity for adaptive control and personalized feedback.

Ultimately, this research demonstrates that the convergence of AI, behavioral science, and user-centered design holds transformative potential for the way individuals interact with energy systems. MELISSA exemplifies how intelligent interfaces can democratize access to energy knowledge, promote environmental awareness, and support broader societal goals related to sustainability, equity, and digital inclusion.

Author Contributions: Conceptualization, A.L.M.S., M.G.M.P. and G.P.R.F.; Data curation, G.A.P.R., R.I.M. and G.P.R.F.; Formal analysis, A.L.M.S. and G.D.B.; Investigation, G.A.P.R., R.I.M. and M.G.M.P.; Methodology, A.L.M.S. and M.N.d.O.; Project administration, V.P.G.; Resources, G.A.P.R.; Software, G.D.B.; Supervision, V.P.G. and G.P.R.F.; Validation, V.P.G.; Visualization, M.N.d.O.; Writing—original draft, R.I.M., G.D.B. and M.G.M.P.; Writing—review & editing, M.N.d.O. All authors have read and agreed to the published version of the manuscript.

Funding: This research was funded by Foundation for Research Support of the Federal District (FAP-DF).

Data Availability Statement: The original contributions presented in this study are included in the article. Further inquiries can be directed to the corresponding authors.

Acknowledgments: The authors would like to thank the Foundation for Research Support of the Federal District (FAP-DF) for financially supporting this project.

Conflicts of Interest: The authors declare no conflicts of interest.

References

- Thompson, K. Seasonal Patterns in Residential Energy Use. *Energy Build.* **2023**, *189*, 234–247.
- Serrano, A.L.M.; dos Santos Martins, P.H.; Vergara, G.F.; Bispo, G.D.; Rodrigues, G.A.P.; Mosquera, L.R.; Oliveira, M.N.d.; Neumann, C.; Peixoto, M.G.M.; Gonçalves, V.P. Forecasting Ethanol and Gasoline Consumption in Brazil: Advanced Temporal Models for Sustainable Energy-Management. *Energies* **2025**, *18*, 1501. [\[CrossRef\]](#)
- International Energy Agency. World Energy Outlook 2022. 2022. Available online: <https://www.iea.org/reports/world-energy-outlook-2022> (accessed on 8 July 2025).
- Rodriguez, S.; Martinez, E.; Gonzalez, P.; Sanchez, R. Socioeconomic Factors in Residential Energy Consumption. *Energy Econ.* **2022**, *67*, 355–374. [\[CrossRef\]](#)
- Garcia, M.; Lopez, J.; Fernandez, A.; Ruiz, C. The Impact of IoT on Residential Energy Consumption. *IEEE Trans. Smart Grid* **2022**, *12*, 876–889. [\[CrossRef\]](#)
- Weiser, M. The Computer for the 21st Century. *ACM SIGMOBILE Mob. Comput. Commun. Rev.* **1999**, *3*, 3–11. [\[CrossRef\]](#)
- Abowd, G.; Mynatt, E. Charting Past, Present, and Future Research in Ubiquitous Computing. *ACM Trans.-Comput.-Hum. Interact.* **2000**, *7*, 29–58. [\[CrossRef\]](#)
- Atzori, L.; Iera, A.; Morabito, G. The Internet of Things: A Survey. *Comput. Netw.* **2010**, *54*, 2787–2805. [\[CrossRef\]](#)
- Darby, S. Smart Metering: What Potential for Householder Engagement? *Build. Res. Inf.* **2018**, *46*, 126–139. [\[CrossRef\]](#)
- Faruqui, A.; Sergici, S.; Sharif, A. The Impact of Informational Feedback on Energy Consumption—A Survey of the Experimental Evidence. *Energy* **2010**, *35*, 1598–1608. [\[CrossRef\]](#)
- Rosano-Peña, C.; Pensado-Leglise, M.D.R.; Marques Serrano, A.L.; Bernal-Campos, A.A.; Hernández-Cayetano, M. Agricultural eco-efficiency and climate determinants: Application of dea with bootstrap methods in the tropical montane cloud forests of Puebla, Mexico. *Sustain. Environ.* **2022**, *8*, 2138852. [\[CrossRef\]](#)
- Cook, D.; Augusto, J.; Jakkula, V. Ambient Intelligence: Technologies, Applications, and Opportunities. *Pervasive Mob. Comput.* **2012**, *8*, 18–38. [\[CrossRef\]](#)
- Zhou, K.; Fu, C.; Yang, S. Big Data Driven Smart Energy-Management: From Big Data to Big Insights. *Renew. Sustain. Energy Rev.* **2016**, *56*, 215–225. [\[CrossRef\]](#)
- Chandola, V.; Banerjee, A.; Kumar, V. Anomaly Detection: A Survey. *ACM Comput. Surv.* **2009**, *41*, 1–58. [\[CrossRef\]](#)
- Vazquez-Canteli, J.; Nagy, Z. Reinforcement Learning for Demand Response: A Review of Algorithms and Modeling Techniques. *Appl. Energy* **2019**, *235*, 1072–1089. [\[CrossRef\]](#)
- Mocanu, E.; Mocanu, D.; Nguyen, P.; Liotta, A.; Webber, M.; Gibescu, M.; Sloatweg, J. On-line Building Energy Optimization using Deep Reinforcement Learning. *IEEE Trans. Smart Grid* **2018**, *10*, 3698–3708. [\[CrossRef\]](#)
- Jaradat, M.; Jarrah, M.; Bousselham, A.; Jararweh, Y.; Al-Ayyoub, M. The Internet of Energy: Smart Sensor Networks and Big Data Management for Smart Grid. *Procedia Comput. Sci.* **2015**, *56*, 592–597. [\[CrossRef\]](#)
- Bispo, G.D.; Vergara, G.F.; Saiki, G.M.; Martins, P.H.d.S.; Coelho, J.G.; Rodrigues, G.A.P.; Oliveira, M.N.d.; Mosquera, L.R.; Gonçalves, V.P.; Neumann, C.; et al. Automatic literature mapping selection: Classification of papers on industry productivity. *Appl. Sci.* **2024**, *14*, 3679. [\[CrossRef\]](#)
- Farhangi, H. The Path of the Smart Grid. *IEEE Power Energy Mag.* **2010**, *8*, 18–28. [\[CrossRef\]](#)
- Williams, R. Prosumers in the Energy Transition. *Energy Policy* **2021**, *156*, 112432. [\[CrossRef\]](#)
- Serrano, A.L.M.; Rodrigues, G.A.P.; Martins, P.H.d.S.; Saiki, G.M.; Filho, G.P.R.; Gonçalves, V.P.; Albuquerque, R.d.O. Statistical Comparison of Time Series Models for Forecasting Brazilian Monthly Energy Demand Using Economic, Industrial, and Climatic Exogenous Variables. *Appl. Sci.* **2024**, *14*, 5846. [\[CrossRef\]](#)
- de Oliveira, M.N.; Mosquera, L.R.; Martins, P.H.d.S.; Serrano, A.L.M.; Bispo, G.D.; Vergara, G.F.; Saiki, G.M.; Neumann, C.; Gonçalves, V.P. Tracking Biofuel Innovation: A Graph-Based Analysis of Sustainable Aviation Fuel Patents. *Energies* **2024**, *17*, 3683. [\[CrossRef\]](#)
- Siano, P. Demand Response and Smart Grids—A Survey. *Renew. Sustain. Energy Rev.* **2014**, *30*, 461–478. [\[CrossRef\]](#)
- Tuballa, M.; Abundo, M. A Review of the Development of Smart Grid Technologies. *Renew. Sustain. Energy Rev.* **2016**, *59*, 710–725. [\[CrossRef\]](#)
- Raza, S.; Mokhlis, H.; Arof, H.; Laghari, J.; Wang, L. A Review on Distributed Energy Resources and MicroGrid. *Renew. Sustain. Energy Rev.* **2015**, *56*, 840–855. [\[CrossRef\]](#)
- Rocha Filho, G.P.; Brandão, A.H.; Nobre, R.A.; Meneguette, R.I.; Freitas, H.; Gonçalves, V.P. Host: Towards a low-cost fog solution via smart objects to deal with the heterogeneity of data in a residential environment. *Sensors* **2022**, *22*, 6257. [\[CrossRef\]](#) [\[PubMed\]](#)
- Chen, H.; Li, D. Natural Language Interfaces for Energy-Management. *Appl. Energy* **2022**, *278*, 115723.
- Brown, T.; Mann, B.; Ryder, N.; Subbiah, M.; Kaplan, J.; Dhariwal, P.; Neelakantan, A.; Shyam, P.; Sastry, G.; Askell, A.; et al. Language Models are Few-Shot Learners. *Adv. Neural Inf. Process. Syst.* **2020**, *33*, 1877–1901.
- Bommasani, R.; Hudson, D.; Adeli, E.; Altman, R.; Arora, S.; von Arx, S.; Bernstein, M.; Bohg, J.; Bosselut, A.; Brunskill, E.; et al. On the Opportunities and Risks of Foundation Models. *arXiv* **2021**, arXiv:2108.07258. [\[CrossRef\]](#)

30. Ouyang, L.; Wu, J.; Jiang, X.; Almeida, D.; Wainwright, C.; Mishkin, P.; Zhang, C.; Agarwal, S.; Slama, K.; Ray, A.; et al. Training Language Models to Follow Instructions with Human Feedback. *Adv. Neural Inf. Process. Syst.* **2022**, *35*, 27730–27744.
31. Raffel, C.; Shazeer, N.; Roberts, A.; Lee, K.; Narang, S.; Matena, M.; Zhou, Y.; Li, W.; Liu, P. Exploring the Limits of Transfer Learning with a Unified Text-to-Text Transformer. *J. Mach. Learn. Res.* **2020**, *21*, 1–67.
32. Radford, A.; Kim, J.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. Learning Transferable Visual Models From Natural Language Supervision. In Proceedings of the 38th International Conference on Machine Learning, Virtual, 18–24 July 2021; Volume 139, pp. 8748–8763.
33. Rivkin, D.; Hogan, F.; Feriani, A.; Konar, A.; Sigal, A.; Liu, X.; Dudek, G. AIoT Smart Home via Autonomous LLM Agents. *IEEE Internet Things J.* **2025**, *12*, 2458–2472. [[CrossRef](#)]
34. Paneru, C.P.; Tarigan, A.K.; Toşa, C. Exploring user engagement and perceptions towards smart energy applications in Norway: Insights from a text mining analysis of user reviews. *Energy Effic.* **2025**, *18*, 25. [[CrossRef](#)]
35. Hosseini, S.; Kelouwani, S.; Agbossou, K.; Cardenas, A.; Henao, N. A semi-synthetic dataset development tool for household energy consumption analysis. In Proceedings of the 2017 IEEE International Conference on Industrial Technology (ICIT), Toronto, ON, Canada, 22–25 March 2017; IEEE: New York, NY, USA, 2017; pp. 564–569.
36. Badar, A.Q.; Anvari-Moghaddam, A. Smart home energy-management system—a review. *Adv. Build. Energy Res.* **2022**, *16*, 118–143. [[CrossRef](#)]
37. Cascone, L.; Sadiq, S.; Ullah, S.; Mirjalili, S.; Siddiqui, H.U.R.; Umer, M. Predicting household electric power consumption using multi-step time series with convolutional LSTM. *Big Data Res.* **2023**, *31*, 100360. [[CrossRef](#)]
38. Johnson, M.R.; Wilson, J. The Evolution of Home Energy-Management Systems: A Historical Perspective. *J. Sustain. Energy Technol.* **2023**, *45*, 112–128.
39. Morales, D.L.; Santana, E.J.; Pereira, A.C. Generational Development of HEMS: From Monitoring to Intelligence. *Energy Inform. Rev.* **2022**, *18*, 203–219.
40. International Energy Agency. *Residential Energy Efficiency Potential: Global Assessment Report*; IEA Publications: Paris, France, 2024.
41. Zhang, L.; Bhuiyan, M.Z.A.; Li, J. Integrating Large Language Models with Internet of Things: Applications, Challenges, and Future Directions. *IEEE Internet Things J.* **2024**, *11*, 8721–8739. [[CrossRef](#)]
42. Kumar, S.; Rodrigues, J.J.P.C. Federated Edge Computing for IoT-LLM Integration in Energy Systems. *Sustain. Comput. Inform. Syst.* **2023**, *37*, 100789.
43. Almeida, T.R.; Sousa, P.M. Energy Efficiency of Edge vs. Cloud Processing in IoT-LLM Architectures. *J. Green Comput.* **2024**, *9*, 145–162.
44. Carvalho, M.S.; Oliveira, R.A. Optimizing DBSCAN Parameters for Anomaly Detection in Residential Energy Data. *Energy Build.* **2023**, *278*, 112673.
45. Ferreira, L.N.; Santos, C.M. Dimensionality Reduction Techniques for IoT Energy Data: A Comparative Analysis. *Appl. Energy* **2024**, *345*, 121034. [[CrossRef](#)]
46. Ribeiro, I.M.; Providência, B. *Design Centrado no Usuário: Concepções, Práticas e Soluções*; Universidade do Minho: Braga, Portugal, 2021.
47. Martins, P.L.; Costa, R.S. Socioeconomic Disparities in HEMS Adoption: Evidence from Latin American Countries. *Energy Res. Soc. Sci.* **2023**, *96*, 102952. [[CrossRef](#)]
48. Silva, J.C.; Pereira, M.A. Privacy Concerns in Residential Energy Monitoring: Challenges and Solutions. *IEEE Trans. Smart Grid* **2024**, *15*, 2187–2201.
49. Rodrigues, A.F.; Mendes, T.V. Multi-Agent Architectures for Enhanced Privacy in Smart Energy Systems. *J. Inf. Secur. Appl.* **2023**, *74*, 103387.
50. Rosano-Pena, C.; De Almeida, C.A.R.; Rodrigues, E.C.C.; Serrano, A.L.M. Spatial dependency of eco-efficiency of agriculture in São Paulo. *Braz. Bus. Rev.* **2020**, *17*, 328–343. [[CrossRef](#)]
51. Sustainable Energy Research Group. *Carbon Footprint Reduction Potential of Smart Home Technologies*; Technical report; University of Southampton: Southampton, UK, 2025.
52. Oliveira, C.M.; Fernandes, E.S. Life Cycle Assessment of Smart Home Energy-Management Systems. *J. Clean. Prod.* **2024**, *425*, 138721.

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.