MELISSA: An LLM-Powered Smart Home Energy Consumption Monitoring Framework

Gabriel Arquelau Pimenta Rodrigues[1], **Matheus Noschang de Oliveira**[1], **André Luiz Marques Serrano**[1], **Geraldo Pereira Rocha Filho**[1,2], **Guilherme Fay Vergara**[1], **Letícia Rezende Mosquéra**[1], **Vinícius Pereira Gonçalves**[1]

[1]Electrical Engineering Department (ENE) – University of Brasilia (UnB)
Brasilia – Brazil

[2]Department of Exact and Technological Sciences (DCET)
State University of Southwest Bahia (UESB) – Vitória da Conquista – Brazil

{gabriel.arquelau,matheus.oliveira,guilherme.vergara,
leticia.mosquera}@redes.unb.br, {andrelms,vpgvinicius}@unb.br,
geraldo.rocha@uesb.edu.br

Abstract. *This work introduces MELISSA, a multi-agent system that uses Large Language Models (LLMs) to optimize household energy consumption by integrating historical analysis and meteorological inputs, acting as an intelligent Home Energy Management System (HEMS) for smart spaces. Through a two-stage process that condenses data by approximately 99%, the system identifies consumption patterns and anomalies. The Gunning Fog Index indicates that the outputs are easily readable by the target audience, with a moderate Self-BLEU score. Thus, MELISSA offers an effective residential energy management solution, using LLMs to communicate with the end-users. Future enhancements include integrating energy generation data.*

1. Introduction

Energy management systems monitor and control energy resources to enhance efficiency and sustainability. This approach encompasses measuring energy consumption, implementing strategies to optimize electricity use, and promoting the adoption of renewable energy sources. When applied effectively, these systems minimize environmental impacts while significantly reducing operational and household costs. As residential energy use represents a significant portion of global electricity demand [Wang et al. 2024], these benefits emphasize the importance of intelligent systems in optimizing household energy consumption.

The rise of urbanization and increasing reliance on electronic devices have intensified energy use in residential environments, making it essential to adopt more efficient practices. However, many households still struggle with inefficient consumption patterns, whether due to improper appliance use, lack of control over connected devices or the absence of tools for intelligent energy management.

With the growing adoption of ubiquitous and pervasive computing, an emerging opportunity exists to transform how energy is used in residential settings. Automated systems enable more efficient device management, whether through remote control of appliances, automatic consumption adjustments based on user needs or integration with renewable energy sources such as solar panels. The connectivity provided by the Internet of Things (IoT) has promoted this evolution, allowing devices to communicate with each other and respond to real-time inputs, optimizing electricity usage and reducing waste.

In this context, the Internet of Energy (IoE) is a solution to improve energy efficiency and solve uncontrolled consumption. The IoE refers to the interconnection of intelligent energy devices within a digital ecosystem that enables consumption monitoring and automation. Integrating smart home technologies with advanced energy management systems facilitates the modernization of traditional power grids, transforming them into intelligent networks capable of distributing energy more efficiently and sustainably.

Integrating the IoE and Large Language Models (LLMs) represents a significant advancement in energy management. These models can transform how energy data is interpreted, making information more accessible for users. The LLMs can process data generated by smart home devices and provide personalized recommendations based on individual consumption habits. Instead of users having to interpret numerical values about their energy consumption, they can interact with an LLM-based virtual assistant that offers detailed guidance on reducing waste and improving energy efficiency [Giudici et al. 2024] with a more user-friendly interface that enables individuals to interpret its results. Furthermore, LLMs facilitate more natural communication between consumers and energy systems, allowing real-time responses to inquiries and personalized strategy suggestions.

Based on these technological advancements, we introduce MELISSA (Modern Energy LLM-IoE Smart Solution for Automation), a smart Home Energy Management System (HEMS) solution. MELISSA is designed to advance the interaction between individuals and their household energy consumption information, offering a novel approach to optimizing electricity demand. MELISSA continuously monitors a household's energy consumption, using IoE integration to collect electricity usage data. Analyzing consumption patterns provides personalized recommendations to reduce waste and alerts users when household appliances may require maintenance. Additionally, MELISSA communicates with users using natural language, facilitating data comprehension.

The key benefits of MELISSA include reducing energy waste by automatically identifying devices consuming excessive electricity, enabling financial savings by helping users adopt cost-cutting strategies, promoting sustainability by encouraging more conscious energy use, and offering ease of use through an intuitive AI-powered interface that allows personalized communication.

2. Literature review

Several studies have explored the development of a HEMS to optimize household energy consumption [Badar and Anvari-Moghaddam 2022]. One such approach focuses on generating long-term semi-synthetic datasets to address the scarcity of real-world energy consumption data [Hosseini et al. 2017].

By simulating a simple household environment and employing probabilistic models based on statistical analysis of real-world data, the authors provide stochastic power profiles that enhance the accuracy of energy management evaluations. In alignment with this research, MELISSA uses data-driven knowledge to enhance user interaction with household energy systems, integrating intelligent analysis and recommendations to improve energy efficiency and sustainability.

Advancing these efforts to generate data, various studies have provided real-world energy consumption datasets to support further research. As energy consumption may vary according to regional factors, such as climate, these datasets are specific to their respective countries. Examples of datasets include those from Chile [Condon et al. 2022], and from the UK [Pullinger et al. 2021]. These datasets may be used, for example, for forecasting energy consumption and for detecting anomalies [Liu et al. 2021].

Beyond the energy sector, LLM has been used for smart home automation. For instance, Sasha [King et al. 2024] interprets vague requests (e.g. ”help me see better”) to create automation routines (e.g. turning on the lights). This approach can be applied to HEMS with LLMs, enabling intuitive energy management by interpreting vague user commands to analyze consumption. The work of [Michelon et al. 2025] also focuses on using LLMs to format data into well-formatted preferences for HEMSs systems to work better. Similarly, LLMs have been explored for smart home simulation, where they generate human-like daily use activities for virtual agents, reducing the complexity of user configurations and increasing the adaptability of automation routines [Yonekura et al. 2024]. This adaptability can be leveraged in HEMS, where LLMs could generate personalized energy-saving routines based on learned user behavior and environmental conditions.

3. MELISSA architecture

This solution uses an LLM to continuously assess energy data from a household, considering its historical use, local weather data and number of occupants in the house. The steps of the architecture used to enable this Human-Computer Interaction (HCI) are depicted in Figure 1, which encompasses the data collection, as more deeply explained in Section 3.1; and its integration with the LLM API, as detailed in Section 3.2. All the experiments were conducted using Python version 3.11.11.

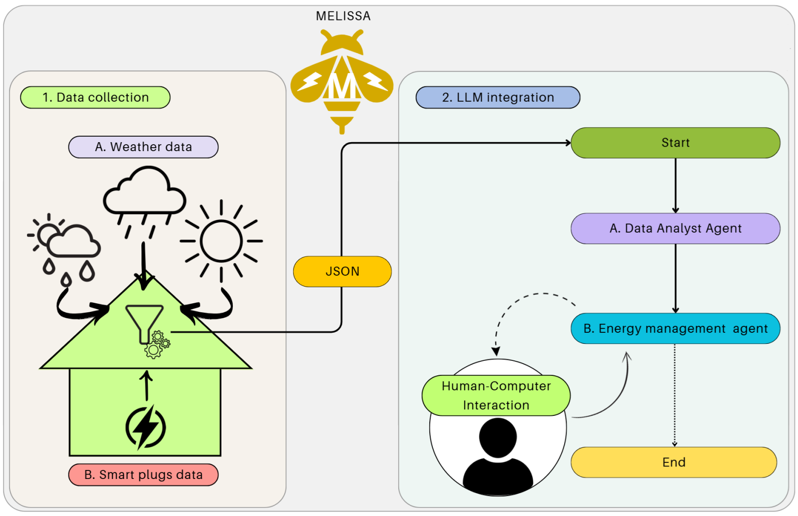

Figure 1. MELISSA framework

3.1. Data collection

The first step in analyzing and monitoring household energy consumption is data collection. All collected data is structured in JavaScript Object Notation (JSON) format and recorded at an hourly frequency. It is later inputted into the LLM autonomous agents.

For energy consumption measurement, we use TP-Link's Tapo P110 smart plugs, as they may be remotely accessed and controlled through an API[1]. These devices have built-in energy monitoring features that provide data on energy consumption, measured in kilowatt-hours (kWh). The data spans from February 12, 2025, to March 14, 2025.

In the experimental setup, the smart plugs are installed to measure the consumption of a refrigerator, a water dispenser, a microwave, a fan and a washing machine in a two-person household in Brasília, Brazil. MELISSA tracks the electricity usage within the household and transmits this data to a central system for analysis and monitoring. These common household appliances were selected to represent typical residential energy profiles.

In addition to that, meteorological data is relevant in an HEMS context because weather conditions directly impact household energy usage [Serrano et al. 2024]. For example, temperature fluctuations influence the demand for cooling systems. Furthermore, incorporating meteorological data allows for a more accurate energy efficiency analysis, providing knowledge of how external conditions influence consumption. MELISSA collects weather data using the Open-Meteo API, the air temperature measured 2 meters above the ground, the apparent temperature (both in °C), the wind speed measured 10 meters above the ground (in km/h), the total cloud cover (in %) and the precipitation (in mm).

3.2. LLM integration

The JSON data generated during collection is fed into the LLM API for analysis. While we currently use Gemini for this purpose, the system is designed to be adaptable. It can easily be configured to work with other preferred LLM models, such as DeepSeek, Claude, or GPT, enabling analysis based on the chosen model's specific capabilities.

MELISSA employs two autonomous agents: the data analyst agent, which processes the collected JSON to summarize the raw inputted data; and the energy management agent, which interprets the data provided by the data analyst agent, along with last week's raw data to provide knowledge and interact with the user. These agents are designed to work synergistically, with the data analyst agent focusing on data-driven analysis and the energy management agent taking action based on the study, enabling efficient energy management.

The agents were configured to process with a zero temperature to ensure consistent responses, fewer assumptions, and a maximum output length of 2.000 tokens for the energy management agent, to make the response concise. The structured system prompts were developed using the RICES (Role, Instructions, Context, Expectations, Style) [Vogelsang 2024], combined with the CLEAR frameworks [Lo 2023], to delineate tasks for each AI agent. This approach maximizes their potential for accurate task execution. These agents summarize energy consumption patterns and help identify trends that

[1]https://github.com/mihai-dinculescu/tapo

may suggest opportunities for improving energy efficiency or indicate areas where consumption may be unnecessarily high. For instance, they might detect that energy usage spikes during certain hours or correlate it with specific appliances, allowing the user to ask further questions about their energy consumption. In addition, the analysis considers possible anomalies, which could indicate issues such as malfunctioning appliances or unanticipated energy surges.

3.2.1. Data analyst agent

The data analyst agent is designed to process and summarize the hourly datasets related to energy consumption and meteorological variables. The system prompt created for this agent is depicted in Table 1. This agent is tasked with analyzing historical energy consumption data for household appliances while considering critical contextual factors, such as the number of tenants in the household, which can influence energy usage patterns.

The data analyst agent serves as the first step in interpreting the collected data, and its output is fundamental for providing interpretable information for the energy management agent. The prompt is structured to ensure that the agent focuses on delivering a comprehensive yet concise summary of the dataset, ensuring it is usable for the next step in the workflow. The following quote contains the user prompt given to the agent:

"Summarize the following historical data for a house of $\{num\_tenants\}$ tenants with the maximum possible detail so that the next LLM agent can have a full view of the datasets:$\{combined\_history\_data\}$".

Table 1. Data analyst agent RICES system prompt

Category	Instruction
Role	You are a data analyst with expertise in summarizing data to another LLM. Your role is to provide the minimum information that describes the datasets completely.
Instructions	Analyze the attached historical energy consumption and weather data for household appliances, considering the number of tenants in the household.
Context	The dataset includes historical energy consumption for specified household appliances and weather data for the same date range.
Expectations	Present data in a structured format. Provide sufficient context for another LLM to understand the data completely.
Style	Maintain a professional and concise tone. Avoid jargon unless necessary and explain technical terms when used. Understandably present the findings so that an LLM can conclude.

Token Reduction Ratio (TRR) is calculated as per Equation 1 and is used to evaluate the amount of information that flows onto the next agent compared to its input JSON.

$$\text{TRR} = \left(1 - \frac{T_s}{T_o}\right) \times 100 \tag{1}$$

where T_o is the number of tokens in the original text, and T_s is the number of tokens in the summarized text.

3.2.2. Energy management agent

The system prompt created to feed the energy management agent is depicted in Table 2. This agent is responsible for analyzing the summarized data provided by the data analyst agent and transforming it into information to improve household energy efficiency. This agent examines household appliances' historical energy consumption data while considering contextual factors such as the number of tenants in the household. The output contributes to making informed decisions about optimizing energy use, reducing costs, and minimizing environmental impact.

Serving as the second stage in the data analysis process, the energy management agent complements the findings from the data analyst agent to offer detailed recommendations and enable the HCI regarding energy consumption.

Table 2. Energy management agent RICES system prompt

Category	Instruction
Role	You are an energy data analyst with expertise in historical energy consumption data analysis for household appliances. You provide detailed information and recommendations based on the provided datasets.
Instructions	Analyze the attached historical energy consumption data for household appliances, considering the number of tenants in the household and the weather conditions. Compare the energy usage of each appliance with the most recent states and avoid speculating on anomalies without data evidence.
Context	The dataset includes historical energy consumption data for specified household appliances. The goal is to derive knowledge that can improve energy efficiency, reduce costs, and minimize environmental impact, factoring in the number of tenants in the household and weather conditions.
Expectations	Present information in a structured format. Offer clear explanations. Provide recommendations for optimizing energy usage. Ensure that the analysis is comprehensive and addresses all requested aspects.
Style	Maintain a professional and concise tone. Avoid jargon unless necessary and explain technical terms when used. Present the findings in an understandable way for both technical and non-technical stakeholders.

After passing the system prompts to the respective agents, a function to analyze energy efficiency was created with the following command to the energy management agent, in which $summary$ the output generated from the previous agent, $recent\_data$ corresponds to last week's energy data, and $recent\_weather\_data$ represents the previous week's weather data. Next is the user prompt that was inputted to the second agent:

"Based on this summary ($\{summary\}$), compare it with recent data ($\{recent\_data\}$) and weather conditions ($\{recent\_weather\_data\}$). Suggest cost-saving measures and alert if any anomaly is detected."

The Gunning Fog Index, as defined in Equation 2, is used to evaluate the outputs of the energy management agent. It is a readability metric that estimates the years of formal education a person needs to understand a piece of text on first reading. A higher score indicates that the text is more complex.

$$\text{GFI} = 0.4 \times \left(\frac{\text{Words}}{\text{Sentences}} + 100 \times \frac{\text{Complex Words}}{\text{Words}} \right) \tag{2}$$

We also use Self-BLEU, a diversity metric that measures the similarity among generated texts, defined as the average BLEU score of each text against all others in the set. It ranges from 0 to 1, with higher values indicating similar texts and lower values indicating higher diversity.

3.2.3. Human-computer interaction

As the primary users of these agents are typically non-experts in data science, ensuring that the interaction is user-friendly is critical to the success of these tools. The interaction with MELISSA should be designed to provide relevant summaries that are easy to understand. The user must quickly gain key knowledge, such as energy consumption trends and anomalies, leading to data-driven decision-making. Furthermore, the interaction design should prioritize minimizing cognitive overload, where the user is not overwhelmed by too much detail but can access new information as needed.

A significant advantage of HCI in this context is the ability for users to interact with the system using natural, free-form questions. This allows users to inquire about their energy consumption and receive recommendations without being restricted to predefined queries or specific scopes, making MELISSA more flexible and accessible. This feature enables users to explore various aspects of their energy usage on demand, ask about particular appliances or seek general advice on improving energy efficiency.

4. Results and discussion

This section presents the results from this research to evaluate the `gemini-2.0-flash-thinking-exp-01-21` model's performance in sufficiently describing the data to itself so that it can compare it with recent weekly data without losing information. Furthermore, the output from the second agent is presented as the main result and conclusions are drawn from its performance on both tasks.

4.1. Data Analyst Agent's output and assessment

The first agent is responsible for summarizing the original dataset while preserving critical information. Its output successfully extracts information from the energy consumption and weather datasets, which can be accessed at a public repository [2]. The summarization

[2]zenodo.org/records/15034634

process effectively captured energy consumption data by maintaining hourly granularity, appliance-level differentiation, and timestamps.

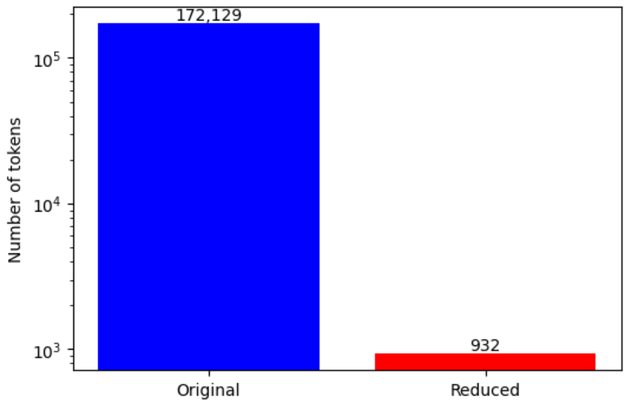

Figure 2. Token reduction by the data analyst agent

The agent reduced the token count from 172,129 to 932, achieving a significant TRR of 99.45%, as depicted in Figure 2. Nevertheless, the output retained the details necessary for further analysis, demonstrating an effective balance between conciseness and informativeness, reducing computational costs while maintaining analytical value.

4.2. Energy Analyst Agent's output and assessment

The second agent analyzed the recent energy consumption dataset compared to the summarized historical data from a week before the last recorded date. The resulting output detailed each appliance's energy usage and recommended cost-saving measures. The analysis effectively captured appliance-level consumption patterns. The refrigerator showed stable energy use with minor fluctuations from defrost cycles or door openings. The water dispenser had low, intermittent consumption. Fan usage correlated with temperature variations, while the washing machine exhibited infrequent use with energy spikes during washing cycles. The microwave displayed short bursts of high consumption, consistent with its operation. Additionally, no significant anomalies were detected in the energy consumption trends, and all observed variations were within expected ranges.

The agent suggested optimizing energy usage by maintaining refrigerator door seals and setting optimal temperatures, minimizing unnecessary water dispenser cooling, strategically using fans with timers and natural ventilation, optimizing washing loads for efficiency, and promoting efficient microwave use while reducing standby power consumption. The recommendations were practical and well-aligned with the observed consumption data.

To further assess the agent's output, we have conducted a readability evaluation on 30 different messages using the Gunning Fog score. Figure 3 shows the results, with a 10.08 ± 0.75 average score, suggesting the outputs could be easily read by 8th to 11th grade students [Gunning 1969]. This finding reinforces that MELISSA's target audience — homeowners — are unlikely to face language barriers when interacting with the system.

Additionally, these 30 outputs achieved a Self-BLEU score of 0.53, reflecting moderate diversity among the generated messages. This level of diversity can be advanta-

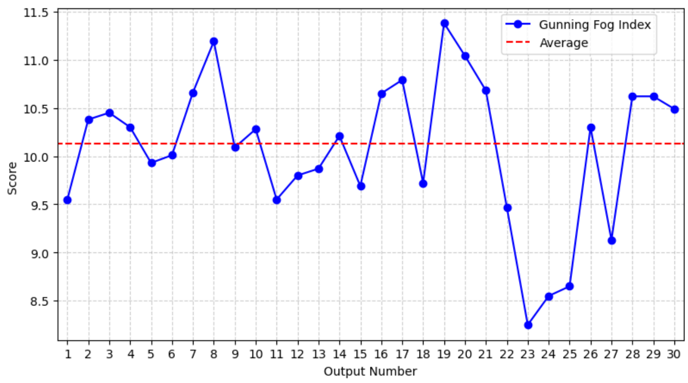

Figure 3. Gunning Fog readability scores

geous, as it ensures that users receive varied responses to avoid monotony yet consistently sufficient to preserve the alignment with the system's communication goals.

5. Conclusions and future works

We propose MELISSA, a framework designed to monitor a household's energy consumption, compare it with the last week's data, and generate recommendations to optimize energy usage, reduce environmental impact, and detect anomalies that could trigger maintenance actions.

The first agent successfully condensed the original dataset by roughly 99%, significantly reducing the computational burden on the second agent, the energy analyst. Despite this compression, the energy analyst effectively extracted meaningful information, accurately identified consumption patterns, correlated energy use with external factors, and provided personalized recommendations. This demonstrates the efficiency and reliability of MELISSA's multi-agent architecture in synthesizing energy datasets and in enabling human interaction with household energy consumption patterns in a readable format.

Future works could explore the integration with actuators, enabling direct control over household appliances to implement optimizations autonomously. Also, energy generation data could be inputted into MELISSA, along with consumption. Furthermore, incorporating quantitative estimates of potential cost reductions for each recommendation would provide a more robust overview of practical benefits. Although no significant anomalies were identified, employing statistical or machine learning methods for outlier detection, such as z-score calculations or anomaly clustering, would enhance the reliability of the evaluation by identifying any unexpected spikes or abrupt drops in consumption. This could be tested in simulated scenarios, such as maintaining an open refrigerator door for extended durations. MELISSA could also be tested in different physical environments, such as universities.

To broaden the scope of the analysis, a deeper examination of the correlation between energy consumption and weather data is proposed. This can include predictive modeling to determine how additional weather variables influence overall consumption.

Acknowledgments

The authors would like to thank the support of FAP-DF and of the University of Brasilia.

References

Badar, A. Q. and Anvari-Moghaddam, A. (2022). Smart home energy management system–a review. *Advances in Building Energy Research*, 16(1):118–143.

Condon, F., Martínez, J. M., Eltamaly, A. M., Kim, Y.-C., and Ahmed, M. A. (2022). Design and implementation of a cloud-iot-based home energy management system. *Sensors*, 23(1):176.

Giudici, M., Padalino, L., Paolino, G., Paratici, I., Pascu, A. I., and Garzotto, F. (2024). Designing home automation routines using an llm-based chatbot. *Designs*, 8(3):43.

Gunning, R. (1969). The fog index after twenty years. *Journal of Business Communication*, 6(2):3–13.

Hosseini, S., Kelouwani, S., Agbossou, K., Cardenas, A., and Henao, N. (2017). A semi-synthetic dataset development tool for household energy consumption analysis. In *2017 IEEE International Conference on Industrial Technology (ICIT)*, pages 564–569. IEEE.

King, E., Yu, H., Lee, S., and Julien, C. (2024). Sasha: Creative goal-oriented reasoning in smart homes with large language models. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.*, 8(1).

Liu, X., Ding, Y., Tang, H., and Xiao, F. (2021). A data mining-based framework for the identification of daily electricity usage patterns and anomaly detection in building electricity consumption data. *Energy and Buildings*, 231:110601.

Lo, L. S. (2023). The clear path: A framework for enhancing information literacy through prompt engineering. *The Journal of Academic Librarianship*, 49(4):102720.

Michelon, F., Zhou, Y., and Morstyn, T. (2025). Large language model interface for home energy management systems. *arXiv preprint arXiv:2501.07919*.

Pullinger, M., Kilgour, J., Goddard, N., Berliner, N., Webb, L., Dzikovska, M., Lovell, H., Mann, J., Sutton, C., Webb, J., et al. (2021). The ideal household energy dataset, electricity, gas, contextual sensor data and survey data for 255 uk homes. *Scientific Data*, 8(1):146.

Serrano, A. L. M., Rodrigues, G. A. P., Martins, P. H. d. S., Saiki, G. M., Filho, G. P. R., Gonçalves, V. P., and Albuquerque, R. d. O. (2024). Statistical comparison of time series models for forecasting brazilian monthly energy demand using economic, industrial, and climatic exogenous variables. *Applied Sciences*, 14(13):5846.

Vogelsang, A. (2024). From specifications to prompts: On the future of generative large language models in requirements engineering. *IEEE Software*, 41(5):9–13.

Wang, T., Zhao, Q., Gao, W., and He, X. (2024). Research on energy consumption in household sector: a comprehensive review based on bibliometric analysis. *Frontiers in Energy Research*, 11.

Yonekura, H., Tanaka, F., Mizumoto, T., and Yamaguchi, H. (2024). Generating human daily activities with llm for smart home simulator agents. In *2024 International Conference on Intelligent Environments (IE)*, pages 93–96.