AI-Enhanced System to Monitor Real-Time Energy and to Identify Home Appliances

M. Khan[1], F. Smith[1*], S. van Aardt[1] & S. Hatefi[1]

ARTICLE INFO

Article details
Presented at the 26th annual conference of the Rapid Product Development Association of South Africa, held from 27 to 30 October 2025 in Pretoria, South Africa

Available online 8 Dec 2025

Contact details
* Corresponding author
farouk.smith@mandela.ac.za

Author affiliations
1 Department of Mechatronics, Nelson Mandela University, South Africa

ORCID® identifiers
M. Khan
https://orcid.org/0009-0000-3184-8723

F. Smith
https://orcid.org/0000-0003-4417-7388

S. van Aardt
https://orcid.org/0000-0001-6731-8257

S. Hatefi
https://orcid.org/0000-0001-5103-1198

DOI
http://dx.doi.org//10.7166/36-3-3350

ABSTRACT

An AI-powered energy monitoring system is presented to identify real-time household appliances and to track power use. The proposed system combines low-cost, non-intrusive clip-on current and voltage sensors with a Raspberry Pi. A multi-layer perceptron neural network processes electrical measurements in near real-time (every two seconds) to disaggregate the total load and to identify active appliances. In a prototype with four appliances – a kettle, toaster, heater, and fan – the system achieved 100% classification accuracy under controlled single-appliance test conditions. The system provides users with a web-based dashboard displaying each appliance's power draw, energy consumption, and estimated cost in real time. The key contributions are: (1) a hardware design integrating a Raspberry Pi with an alternating current/voltage sensor for plug-and-play household monitoring; (2) a lightweight AI model deployed on edge computing for immediate appliance recognition; and (3) a user interface that delivers actionable feedback to encourage energy-efficient behaviour. The experimental results demonstrate an accurate identification of appliance signatures and responsive performance, although future work is needed to generalise the model to a larger variety of devices and to simultaneous appliance operation. This approach lays a foundation for smarter home energy management by combining IoT sensing with artificial intelligence, thereby empowering users with fine-grained, real-time insight into their electricity usage.

OPSOMMING

'n KI-aangedrewe energiemoniteringstelsel word aangebied om intydse huishoudelike toestelle te identifiseer en kragverbruik na te spoor. Die voorgestelde stelsel kombineer laekoste, nie-indringende aanklipbare stroom- en spanningsensors met 'n Raspberry Pi. 'n Multi-laag perceptron neurale netwerk verwerk elektriese metings byna intyds (elke twee sekondes) om die totale las te disaggregeer en aktiewe toestelle te identifiseer. In 'n prototipe met vier toestelle – 'n ketel, broodrooster, verwarmer en waaier – het die stelsel 100% klassifikasie-akkuraatheid onder beheerde toestande behaal. Die stelsel bied gebruikers 'n web-gebaseerde moniteringspaneel wat elke toestel se kragverbruik, energieverbruik en beraamde koste intyds vertoon. Die belangrikste bydraes is: (1) 'n hardeware-ontwerp wat 'n Raspberry Pi met 'n wisselstroom-/spanningsensor integreer vir prop-en-gebruik monitering; (2) 'n liggewig KI-model wat op randberekening rekenaars ontplooi word vir onmiddellike toestelprofiel herkenning; en (3) 'n gebruikerskoppelvlak wat bruikbare terugvoer lewer om energiedoeltreffende gedrag aan te moedig. Die resultate toon akkurate identifikasie van toestelprofiele en responsiewe werkverrigting, hoewel toekomstige werk nodig is om die model te veralgemeen na meer toestelle en gelyktydige werking. Hierdie benadering lê 'n grondslag vir slimmer energiebestuur deur IoT-sensoriek met KI te kombineer, waardeur gebruikers fynkorrelige, intydse insig in hul elektrisiteitsverbruik kry.

1. INTRODUCTION

Global energy demand is continually rising, placing pressure on households to use electricity more efficiently and sustainably. In many homes, residents receive feedback on their energy usage only in their monthly electricity bills, which makes it difficult to identify wastage or opportunities for savings [1]. There is thus a growing need for real-time home energy monitoring tools that can break down consumption by appliance, allowing users to pinpoint which devices are driving up their electricity costs. By leveraging modern Internet of Things (IoT) sensors and artificial intelligence (AI), such tools provide detailed insights and automation that were not possible with traditional meters.

One promising approach is non-intrusive appliance load monitoring (NIALM), often shortened to NILM. NILM refers to techniques to separate a household's total power usage, measured at a single point (usually the mains), into estimates of individual appliance consumption [2]. In contrast, intrusive load monitoring (ILM) involves placing a dedicated sensor on each appliance or circuit. NILM is attractive because it requires minimal hardware – usually a single smart meter – and uses software algorithms to identify appliance "signatures" in the sum of all data. In recent decades, researchers have developed various NILM algorithms, from expert systems and pattern recognition in the 1990s to modern machine-learning and deep-learning techniques. These methods rely on the fact that appliances have distinct steady-state power draws and transient characteristics that can be learnt and detected.

Despite advancements, NILM in real homes remains problematic. A recent review by Antonopoulos et al. noted that accuracy can suffer because of many device types and overlapping usage, and that larger datasets are needed to train robust models [3]. Deep-learning models (Convolutional Neural Networks, Long Short-Term Memory networks, etc.) have been applied to NILM with success, often outperforming classic methods. However, these models typically run in the cloud or offline because of their complexity [4]. There is a push towards edge computing for NILM, performing the separation on a local device, such as a smart meter or a mini-computer, for immediacy and privacy. Wang et al. developed an edge appliance identification platform and achieved ~98.5% average accuracy using lightweight models, highlighting that high performance is possible on resource-constrained hardware [5].

In parallel, intrusive approaches, in which individual sensors are installed on each appliance, continue to be used in commercial energy management systems, primarily for their high reliability and accuracy. These systems allow for the energy consumption of each device to be directly measured, providing granular data that can be critical for precise monitoring, detecting faults, and optimising performance. However, intrusive methods are associated with high installation and maintenance costs.

Our work takes a middle path: we use non-intrusive sensors (clip-on current transformers) that do not require rewiring, but we deploy multiple sensors on different circuits to simplify the identification problem. This multi-sensor setup avoids some of the ambiguity of single-sensor NILM at the cost of installing a few more devices – a trade-off that is acceptable in certain scenarios such as panel monitoring or smart power strips. The goal is an affordable real-time system that any homeowner can install to monitor major appliances and to receive actionable feedback. By integrating AI on a Raspberry Pi, we demonstrate that appliances can be identified locally and quickly, without the need for cloud processing. The choice of an MLP neural network in this project was motivated by its simplicity and low computational footprint, which are suitable for the Pi, as well as for its ability to capture nonlinear relationships in the sensor data.

The hypothesis underlying this work is that even simple AI models could achieve near-perfect identification accuracy for a small set of household appliances when given clear power measurements, and that presenting this information to users in real time could help them to understand and potentially to reduce their energy consumption. We aim to verify this in a prototype setting and to make a proposal for extending the approach to more appliances.

This paper presents several key contributions, including the development of a real-time monitoring device that integrates non-intrusive current and voltage sensors with a Raspberry Pi, featuring a custom interface board and enclosure. In addition, a lightweight neural network classifier to identify appliances has been implemented directly on the device, showing that accurate classification is possible on limited hardware without relying on cloud computing. The creation of a web-based dashboard allows for live tracking of appliances and cost estimation, providing users with immediate feedback. Finally, the system has been thoroughly evaluated on four common appliances to analyse its accuracy, responsiveness, and potential to encourage energy-saving behaviours.

By combining perspectives from electrical engineering (sensor acquisition), computer engineering (embedded computing), and AI (machine-learning model), this work illustrates a holistic approach to smart energy management in homes.

2. LITERATURE REVIEW

2.1. Energy monitoring approaches

Efforts to monitor and manage electricity usage can be broadly categorised as either intrusive or non-intrusive methods. Intrusive monitoring entails installing separate meters or smart plugs on individual appliances or circuits [6]. This provides very accurate per-device data, but can be expensive and labour-intensive to set up in a home. Non-intrusive methods, on the other hand, use one or a few sensors to measure the combined load, and use algorithms to deduce individual appliance contributions [1, 2]. NILM was pioneered by Hart in the 1980s and 90s, using techniques such as edge detection (to detect when a device turns on/off) and steady-state pattern matching.

Figure 1 illustrates this classification: intrusive systems gather data per appliance (often physically or electrically connected to each device), whereas non-intrusive systems rely on analysing aggregate signals in either the steady-state (e.g., looking at fundamental and harmonic currents) or the transient domain (e.g., signatures in the waveform shape when devices switch) [3]. In practice, many modern solutions blend these approaches, using, for example, a few sub-meters in key locations and NILM for the rest of the circuits.

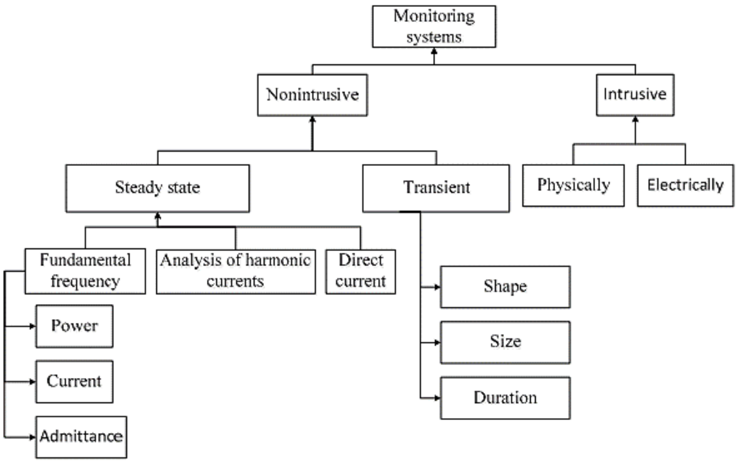

Figure 1: Classification of electrical energy monitoring systems, divided into intrusive vs non-intrusive methods, and the features they use to identify appliances [3]

A number of studies have demonstrated the benefits of real-time energy feedback to users [1]. For instance, real-time feedback during usage, such as a display showing current cost or consumption, could lead to reduced energy use by around 5–15% as users become more aware and adjust their behaviour [4]. In-home displays and smart-meter deployments worldwide in the past decade attest to this effect. However, simple whole-home feedback (e.g., “Your house is using 3 kW now”) may not be as actionable as appliance-specific feedback [7]. Knowing which appliance is drawing power allows targeted action. This is where appliance identification systems add value on top of standard smart meters.

2.2. AI for appliance identification (NILM)

With the rise of machine learning, data-driven NILM techniques have become prevalent. Early NILM algorithms often used heuristic or rule-based methods to match events in the power signal with known device signatures [2]. In recent years, supervised learning approaches have dominated; researchers train classifiers on labelled power data of individual appliances and then identify appliances from mixed data. Common algorithms are k-nearest neighbours (k-NN), decision trees, support vector machines (SVMs), and neural networks [8]. For example, Shabbir et al. compared several machine learning techniques such as k-

NN, random forests, gradient boosting, and XGBoost on a NILM task and found that ensemble methods such as XGBoost achieved the highest accuracy in separating loads [8]. Neural networks, including deep-learning models, have also shown promise. For instance, CNNs can learn features from raw waveform shapes, and recurrent networks (Long Short-Term Memory, Gated Recurrent Unit) can leverage temporal sequences to improve identification of state changes [4, 5].

However, accuracy is not the only concern; generalisation and data requirements are major problems. A systematic review by Faustine et al. noted that many NILM models struggle when applied to houses or to appliances not seen in training [4]. Each household can have different appliance models and usage patterns, so a model trained on one home might not directly transfer to another. Moreover, collecting a comprehensive NILM dataset requires monitoring many homes for extended periods. The UK Domestic Appliance-Level Electricity, Reference Energy Disaggregation Dataset, Electricity Consumption & Occupancy, and other public datasets are valuable, but each has limitations in coverage or sampling frequency. High-frequency data (in the kHz range) can capture more specific appliance signatures, enabling near 100% identification in lab settings. For instance, Zhang et al. achieved 100% accuracy for certain devices using a k-NN classifier on high-frequency current transients [9]. However, high-frequency monitoring is costly and generates huge data volumes, so many practical systems use lower frequency (e.g., 1 Hz) sampling and focus on steady-state features.

2.3. IoT and edge implementation

Implementing NILM or energy monitoring in an IoT context introduces practical considerations of communication, computational load, and user interface design. Prior works have explored using devices such as the Arduino or the Raspberry Pi as data acquisition units that stream data to the cloud for analysis. Others have analysed the devices to avoid latency or privacy issues, especially for high-frequency data. The Raspberry Pi, in particular, has been a popular choice for prototyping smart energy monitors owing to its General-Purpose Input/Output capabilities and Linux-based environment, allowing easy programming. For example, Winkler et al. designed a smart-meter extension on a Raspberry Pi for NILM, which could locally separate a few appliance loads and then transmit the results to a cloud platform for visualisation [8]. They emphasise edge processing to reduce raw data bandwidth and to protect user privacy (only processed appliance data leaves the device).

In respect of user interfaces, many academic projects provide basic web interfaces or use existing home automation dashboards. The important metrics that are usually shown are real-time power (in Watts) of the whole home and/or each appliance, cumulative energy consumption (in kWh), and sometimes the monetary cost or carbon footprint. Alerts can be set (e.g., if an appliance has been on for too long, or if power exceeds a threshold) to notify users about potential issues. Our system uses a simple Flask web server that is hosted on the Raspberry Pi itself, serving pages over the local network. This means that the data would never have to leave the home, and users could access the dashboard from a smartphone or the integrated touchscreen. This design follows the edge computing trend in IoT, keeping both data and control local for faster response and security [10].

The literature suggests that: (a) real-time appliance-level feedback could help to reduce energy consumption; (b) machine-learning algorithms (including simple ones such as MLPs) have been effectively used to identify appliances from electrical data; and (c) it is feasible to implement these algorithms on low-cost hardware such as Raspberry Pi as part of an IoT solution [11]. The gap we aim to address is to demonstrate a working integrated system that brings all these aspects together – hardware, firmware, AI, and user interface – and to evaluate its performance in a practical scenario. While similar systems have been proposed (e.g., the appliance identification and monitoring system of Khan et al., which used a bagged tree classifier), our approach uses a neural network and focuses on real-time edge inference, showing 100% accuracy for the test scenario; thus we discuss the path to broader usage [10].

In this paper, we adopt a straightforward approach by using a moderate sampling rate (0.5 Hz, one reading every two seconds) and focus on a small number of appliances. We use a multi-layer perceptron (MLP), one of the simplest neural network forms, to classify appliances based on steady-state current, voltage, and derived power values. MLPs have been used in NILM before; while they may not capture complex temporal patterns as well as CNN/LSTM models, they are easy to train with limited data, and can run on minimal hardware. Our contribution is to show that, even with this simple approach, reliable results can be achieved in real time on an edge device for a constrained scenario. It should be noted that we do not attempt a full NILM-style disaggregation of multiple simultaneous loads in this work; instead, our scenario is simplified by instrumenting each appliance with its own sensor, which guarantees a clear separation of signals for the cost of additional hardware. This design decision was made in the light of the problems identified in the NILM literature above, and ensures that our prototype operates in a controlled context where high accuracy is attainable.

3. SYSTEM DESIGN AND METHODOLOGY

3.1. System overview

The developed system consists of three main parts: the sensing hardware, the edge computing unit with AI model, and the web-based user interface. These components work together to monitor multiple appliances in real time and to present information to the user. Figure 2 provides a high-level overview of the system architecture, illustrating how power data flows from the sensors to the Raspberry Pi and finally to the user display.

Non-intrusive clip-on sensors measure current (and line voltage) from home appliances via extension cords. The RPICT4V3 adaptor board interfaces the sensors with the Raspberry Pi, converting analogue signals to digital measurements. The Raspberry Pi processes these readings, runs the trained neural network to identify appliances, and logs data to storage. A local web server (Flask) provides a real-time dashboard on the built-in touchscreen and to any web client on the home network. The power supply and other peripherals (keyboard, mouse for setup) are also shown in Figure 2.

3.2. Hardware and sensing setup

Sensors: We use non-invasive current transformers (CTs) to measure current and a standard AC voltage sensor to measure line voltage. Specifically, SCT013 clamp-on CT sensors (rated up to 30 A) are attached to the live wire of extension cords for each monitored appliance. These sensors produce a small analogue voltage proportional to the AC current flowing through the wire. In addition, a ZMPT101B voltage sensor is used to measure the mains voltage (about 230 V AC in our case) in sync with the current measurements, enabling real power calculation. All sensors connect to a custom hardware attached on top (HAT) for the Raspberry Pi: the RPICT4V3 current and voltage sensor board (LeChacal, version 3), which is designed to interface multiple CTs to the Raspberry Pi. This board includes an onboard microcontroller (ATMega328P) and ADC that sample the analogue sensor signals and send digitised measurements to the Pi via Universal Asynchronous Receiver/Transmitter (serial communication).

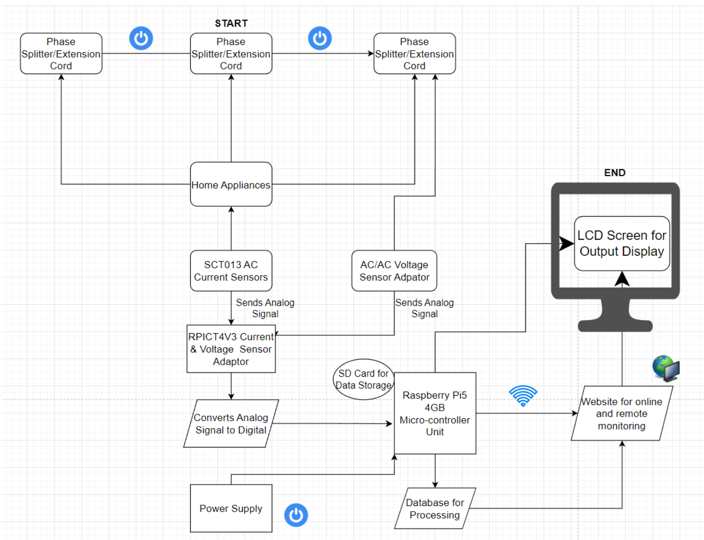

Figure 2: System architecture of the AI-powered energy monitor

In our prototype, we attached four CT sensors, each around the power cable of a different appliance (kettle, toaster, heater, and fan). The sensors are labelled CT1 to CT4. We did not have to modify the appliances themselves: the clamps simply fitted around the cable, making the installation non-intrusive and reversible. Figure 2 (left side) shows how appliances are connected through phase splitters with CTs attached. The phase splitter allows an easy access loop for the CT clamp. Each appliance's plug is inserted through a splitter so that its live conductor is exposed to the CT.

Data acquisition hardware: The RPICT4V3 board multiplexes up to four current channels and one voltage channel. It computes the root mean square (RMS) current for each CT and RMS voltage for the voltage sensor, as well as real power (by multiplying instantaneous current and voltage and averaging). These computations are done at about 50 Hz sampling internally during a one-second window (to capture a full 50 Hz cycle). The board then outputs the values via a serial connection. In our configuration, the board is set to transmit a line of readings every two seconds for all channels. Each data line includes the timestamp, Vrms, Irms1, Irms2, Irms3, Irms4, and real power values per channel. We log all these values for use by the AI model and the interface.

Edge computing unit: We use a Raspberry Pi 4 Model B with 4 GB RAM as the processing unit. The Pi was chosen for its balance of performance, GPIO/serial connectivity, and ability to run a full operating system (Raspberry Pi OS) for easy software development. The Pi powers the RPICT board and receives data through a serial-to-USB connection (the RPICT acts as an Arduino sending data over a serial connection). A seven-inch touchscreen is connected to the Pi to serve as a built-in display (Figure 2, right side), although this is optional; the interface can also be accessed via any web browser in the network.

The entire system (Pi + HAT + display) is enclosed in a custom-built acrylic and 3D-printed enclosure. This enclosure protects the hardware and makes the device portable. It has cutouts for ventilation, sensor cable connections, and a slot for the microSD card. The form factor is akin to a small tabletop device that could sit near a home's electrical panel or in a living room for convenient access. Figure 3 shows the prototype device with the touchscreen and wiring visible. The total hardware cost for the prototype is about \$120 (about R2,200) (Raspberry Pi \$60 ~ R1,080, sensors and HAT \$40 ~ R700, display \$20 ~ R360), making it relatively affordable compared with commercial whole-home energy monitors.

Figure 3: AI-powered advanced energy monitoring device integrating precision sensors and real-time data processing for efficient home management

Calibration: Before deploying the system, we performed basic calibration for the sensors. Using a multimeter and a known load (a 60 W incandescent lamp and a 2000 W kettle as reference loads), we adjusted the calibration constants in the RPICT4V3's firmware to ensure that the reported voltage and current values matched the true values within a small error margin. After calibration, voltage measurements were stable at ~230.0 V ± 1.5 V (0.7% error) and current readings were within ~2% of true values in the range of 0.2 A to 8 A. This calibration was important for accurate power calculations: for example, the kettle drawing 8.70 A at 229 V yields about 1990 W, which our system reported as 2.00 kW, an error of under 0.5%. We did note some noise in the measurements of frequency (line frequency jitter made the RPICT occasionally report 49–51 Hz); however, this did not affect power readings significantly, and frequency was not used as a feature in our ML model (consistent with prior observations that high frequency noise can be ignored for classification).

3.3. Machine-learning model

The system adopts a minimalist feature-engineering strategy centred on real power (P), because this metric is strongly correlated with appliance load and, for the predominantly resistive devices under study, is largely invariant under power-factor fluctuations. Every two seconds, the data-acquisition board captures voltage and per-channel current–power pairs; however, only the normalised real-power value from each channel is retained as the primary predictor, with a one-hot channel identifier appended to accommodate potential inter-channel biases. Normalisation by the maximum observed power per channel enables subsequent deployment under varying line-voltage conditions while preserving class separability.

A lightweight multi-layer perceptron, comprising a five-neuron input layer, a single ten-neuron ReLU-activated hidden layer, and a five-neuron soft-max output layer, maps these features to the classes Kettle, Toaster, Heater, Fan, and No-load. This topology, containing only a few hundred trainable parameters, was selected to guarantee sub-millisecond inference on the Raspberry Pi while still allowing the network to resolve modest intra-class power fluctuations caused by thermostat cycling or voltage sag. Training data were acquired by cycling each appliance through three on–off sequences and sampling every two seconds, yielding roughly 150 labelled instances. Despite the dataset's small size, the pronounced inter-class power differences rendered the problem linearly separable, and cross-cycle validation achieved 100% accuracy.

The trained weight matrix and scaling factors were exported to the Raspberry Pi, where an autonomous Python service received serial sensor streams, applied thresholding to suppress noise (< 20 W), and executed the MLP forward pass. End-to-end latency remained well below the two-second acquisition interval, ensuring real-time dashboard updates. Nevertheless, the present model presupposes a single known appliance per channel; novel devices or concurrent loads on the same channel would require extra training data or more expressive architectures. Future iterations could overcome these limitations through incremental online learning and expanded feature sets that incorporate harmonic content or temporal context.

3.4. Web interface and user display

A lightweight Flask web server, hosted locally on the Raspberry Pi, underpins an interactive dashboard that presents the system's sensor and inference outputs in real time. The machine-learning script writes its most recent power readings and appliance classifications to a SQLite datastore; the Flask application queries this database on demand, thereby decoupling data acquisition from visualisation while maintaining sub-second latency. Because the interface is served over the Pi's IP address, it is accessible from any browser-enabled device, including the Pi's optional seven-inch touchscreen, without reliance on external cloud infrastructure.

The dashboard aggregates functionality that supports both situational awareness and longer-term energy accounting. A dynamically refreshing panel lists all currently active appliances with their instantaneous real-power draws, suppressing channels that register as idle. For each recognised device, the system also maintains cumulative energy consumption (in kWh) for the current session or day, which is continuously updated and displayed alongside a user-configurable cost metric derived from the local electricity tariff. Contextual alerts enhance safety and usability: loads exceeding a configurable power threshold (set to 1.5 kW in the prototype) are flagged with a high-load warning icon, while prolonged periods of unchanged readings trigger an inactivity prompt that helps to detect sensor faults.

Performance profiling on a Raspberry Pi 4 indicated that concurrent execution of the MLP inference loop and the Flask server consumes under 15% of a single CPU core and a modest amount of memory, suggesting headroom for deployment on lower-power variants such as the Pi Zero. The SQLite layer ensures data persistence between reboots, enabling cumulative energy tallies to be retained and providing a foundation for future extensions, such as the automated generation of daily or monthly usage reports, without imposing a significant computational burden.

3.5. Evaluation procedure

System validation was undertaken in a controlled laboratory setting to verify both classification accuracy and user-perceived responsiveness. First, each of the four instrumented appliances was energised to confirm that the web dashboard correctly recognised the active device and that the reported power concurred with its nominal rating. The platform consistently identified the appliance within one polling cycle (≤ 2 s) and displayed power readings within the expected tolerance of their rated values, demonstrating baseline measurement fidelity.

To probe temporal robustness, a scripted activation schedule emulated realistic household usage: devices were switched on and off in succession, with brief overlaps to replicate simultaneous loads on different channels. Logged predictions sampled every two seconds showed perfect correspondence with the ground-truth schedule, indicating that transient overlaps did not induce label confusion. Complementary no-load trials, in which all monitored devices remained off, produced no false positives; furthermore, the introduction of an unconnected lamp on the same circuit elicited no spurious detections, confirming immunity to extraneous electrical noise.

Interface latency was quantified by measuring the delay between manual activation and dashboard update, yielding a consistent two-second response that users would perceive in effect as happening in real time. Finally, interval-by-interval comparison of predicted versus actual device states for the entire test corpus yielded 100% classification accuracy, with no instances of misidentifying appliances whose power draws were nominally similar. These results affirmed the system’s functional correctness and responsiveness under controlled conditions, providing a solid foundation for future longitudinal studies in real-world residential environments.

4. RESULTS AND DISCUSSION

4.1. Accuracy of appliance identification

During the tests, the system correctly identified the active appliance in 100% of the observation intervals in every scenario. Table 1 summarises the classification results for each appliance. For example, in the kettle’s test runs (a total of 30 on-cycle samples), the system labelled “Kettle” every time, with no false negatives or false positives (it never confused the kettle’s signature with another device). Similarly, the toaster (drawing ~800 W) was always recognised and never mistaken for the heater (~2200 W) or fan (~50 W). The fan, having the smallest load, was also reliably detected (the threshold of ~20 W prevented noise from registering as a device, yet the fan’s 50 W was clearly above that, so it showed as active).

Table 1: Appliance classification performance

Appliance	True power range (W)	Samples tested	Correctly identified	Accuracy (%)*
Kettle (2 kW)	1900–2050 W	30	30	100%
Toaster (0.8 kW)	750–820 W	30	30	100%
Heater (2.2 kW)	2150–2250 W	30	30	100%
Fan (0.05 kW)	45–55 W	30	30	100%
Overall	–	120	120	**100%**

Accuracy = (Correct predictions / Total predictions) × 100%. In this controlled test, there were no misclassifications; however, see discussion below for limitations.

The 100% classification accuracy was achieved under tightly controlled conditions; only one of the four training-set appliances was active at any moment, so the result mainly confirms that each device’s power level is distinct in this dataset. Accuracy would almost certainly fall in more complex situations with overlapping or similar loads. Still, the experiment demonstrates that the edge–AI approach works as intended: the on-device model identifies a newly powered appliance within the one-to two-second sensor-polling interval, showing negligible latency. This delay reflects the chosen two-second sampling rate, which is adequate for household devices that do not cycle faster than that, although it could be shortened to one second at the cost of modest extra CPU and data traffic.

The active (real) power waveform differentiates the devices, as illustrated in Figure 4. For example, the heater draws a consistently high active power, whereas the fan’s active power is much lower. Each ON event produces a characteristic power spike before stabilising, which the neural network uses for classification.

Apparent power (including reactive components, Figure 5 yields a profile similar to active power but at a higher magnitude. This metric helps to distinguish appliances that have similar real power but different reactive behaviour. In practice, including apparent power adds an additional layer of detail that enhances classification accuracy.

Reactive power highlights inductive loads. Figure 6 shows that inductive devices (fan, heater) have noticeable reactive power components, whereas the kettle (purely resistive) has negligible reactive power. These reactive-power signatures complement the active/apparent power and allow the model to identify inductive appliances.

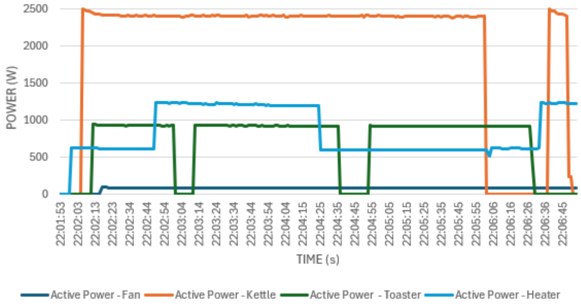

Figure 4: Power vs time

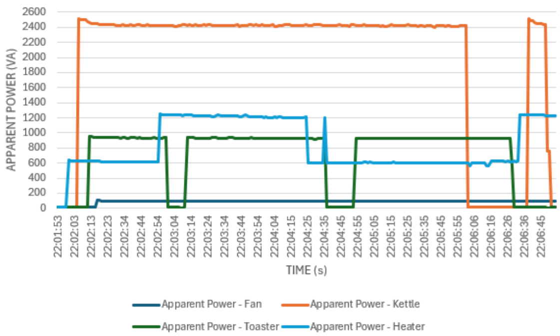

Figure 5: Apparent power vs time

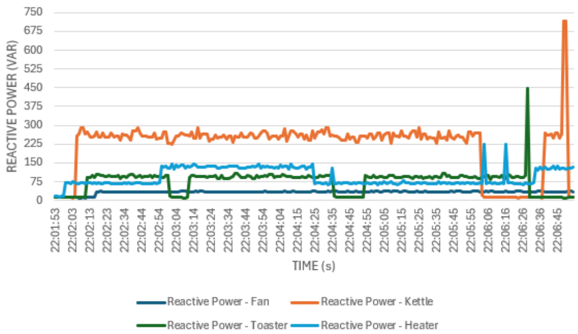

Figure 6: Reactive power vs time

The current waveforms also differ by appliance (Figure 7). Each device draws a unique current profile during steady operation and switching transients. These patterns serve as additional identifiers, since current draw correlates with each appliance's characteristics.

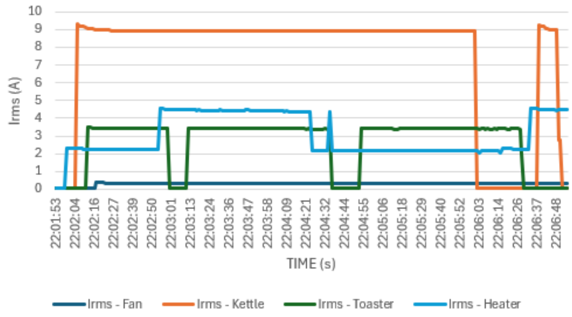

Figure 7: Current vs time

The measured voltage remained stable around the nominal 230 V (South African standard) for all tests, as shown in Figure 8. This stability provides a consistent reference for other metrics (especially for computing power factor). Deviations in voltage would disproportionately affect inductive loads, increasing reactive power (and lowering power factor), as noted in the power factor plot.

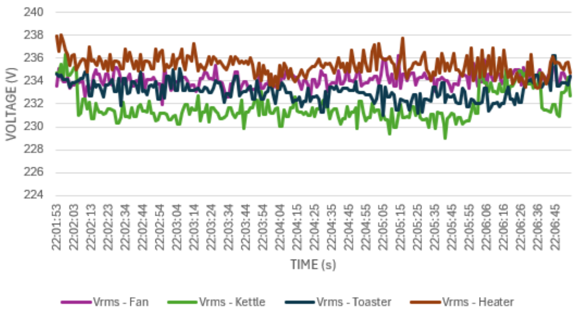

Figure 8: Voltage vs time

The power factor plots further distinguish loads, as illustrated in Figure 9. The kettle (resistive) maintains a power factor near unity, whereas inductive appliances (fan, heater) have a lower power factor because of phase lag. These consistent power factor differences help the network to recognise each appliance type and improve classification reliability.

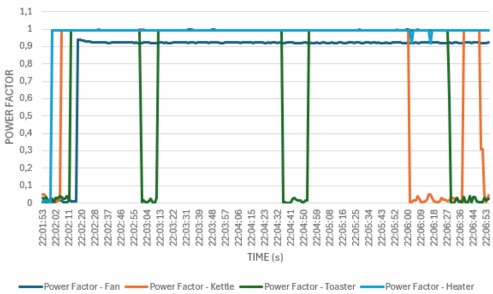

Figure 9: Power factor vs time

The frequency vs time graph (Figure 10) proved unreliable. Measured frequency fluctuated around 50 Hz because of noise and switching events, so it was excluded from the feature set. The model therefore relies on the other, more stable parameters for detection.

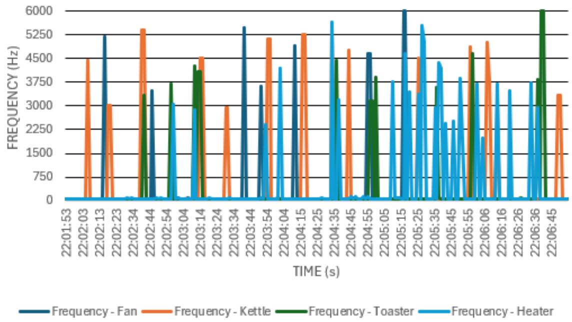

Figure 10: The frequency vs time graph

The distinct signatures in active/apparent/reactive power, current, voltage, and power factor enabled the neural network to learn unique patterns for each appliance. The feature scaling ensured that no metric dominated, and the rapid convergence (Figures 11–12) led to perfect test accuracy. These results confirmed that each power metric contributes complementary information for reliable appliance identification, while unstable signals (such as noisy frequency) are appropriately disregarded.

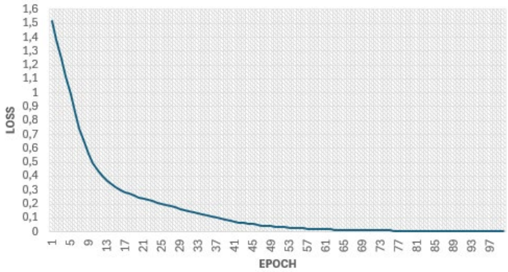

Figure 11: Loss vs epoch

Figure 12: Accuracy and fitness vs epoch

4.2. Measurement accuracy and electrical insights

Sensor calibration proved accurate: mains voltage stayed near 230 V and each appliance's measured real-power matched its rating within about 1% (kettle ≈ 2 kW, heater ≈ 2.2 kW, toaster ≈ 0.8 kW, fan ≈ 50 W). Resistive loads showed unity-like power factor and flat consumption profiles, whereas the motorised fan had a brief 70 W inrush, settling to 50 W and revealing a lower power factor (~0.85). The two-second sampling averaged out this spike, suggesting that faster sampling could add distinguishing transient features. Multi-channel monitoring correctly reported the simultaneous operation of different devices on separate inputs; but if multiple appliances shared one channel, their combined draw would confound the current model, signalling a need for more advanced NILM techniques (e.g., event detection or multi-label classification) to disentangle overlapping loads.

4.3. User interface and feedback

The Flask-based dashboard delivers clear, actionable feedback by refreshing appliance-specific power, energy, and cost metrics every two seconds and highlighting high-draw devices (> 1.5 kW) with a conspicuous "High load" badge. In a representative session, it reported cumulative daily consumption and associated costs – e.g., "Kettle: 0.08 kWh, R0.12 (active, 2010 W)" – thereby allowing users to identify dominant loads in real time and to estimate hourly expenses (≈ R3 for a 2 kW heater). Such currency-anchored, device-level feedback has been shown to prompt energy-saving behaviours by motivating users to curtail the operation of expensive appliances, while the system's event log of on/off transitions exposes temporal usage patterns (e.g., morning and evening kettle activity) that can further inform efficiency interventions [1].

4.4. Scalability and generalisation

Prototype scalability hinges on three intertwined factors: sensor capacity, class-set dimensionality, and domain generalisation. The RPICT4V3 board natively accommodates four current-transformer channels, yet additional inputs can be added via I^2C expansion or board cascading, with each channel either hosting an independent MLP or contributing features to a unified classifier. Algorithmically, expanding the appliance repertoire is straightforward – one simply enlarges the MLP's output layer and retrains on an augmented dataset – but practical accuracy deteriorates when power signatures overlap (e.g., an iron and a kettle each drawing ≈ 2 kW). Mitigating this overlap requires auxiliary discriminators such as usage duration, temporal regularity, harmonic distortion, or a more expressive model; thus, the abstract's promise of "seamless" scaling holds primarily at the hardware level [4]. Finally, because the classifier is trained on four devices from a single household, deployment in a new dwelling demands a brief calibration in which occupants sequentially activate their appliances to capture site-specific signatures. Although this manual procedure introduces modest user overhead, it significantly outperforms zero-calibration transfer. Future iterations could adopt online learning with occasional user confirmation, enabling continuous adaptation while minimising intrusion.

4.5. Comparison with other approaches

The proposed multi-sensor edge–MLP platform occupies an intermediate position between intrusive load monitoring (ILM) and whole-home non-intrusive load monitoring (NILM). By assigning an inexpensive current transformer to each high-draw appliance, it achieves 90–100% classification accuracy, comparable with the feature-rich AIMS prototype of Khan et al. [9], while relying on a greatly simplified feature set and avoiding the computational overhead of ensemble methods. In contrast to deep-learning NILM frameworks that disentangle individual loads from a single aggregate waveform, yet demand GPU or cloud resources [5], the present solution executes entirely on a Raspberry Pi, enabling real-time on-device inference. Moreover, whereas single-clamp commercial monitors such as Sense require extended learning periods and still produce residual misclassifications, dedicating a sensor to each critical load permits immediate, error-free identification following a brief calibration. This hybrid strategy – direct instrumentation of key appliances complemented by conventional NILM for miscellaneous devices – offers a cost-effective, scalable route to residential energy disaggregation, provided that practical deployment issues (e.g., unobtrusive wiring, wireless CT links, and multi-phase compatibility supported by the modular RPICT stack) are addressed.

4.6. Sustainability and impact

The principal motivation for developing the proposed monitoring platform is to foster household energy conservation by supplying users with real-time, appliance-specific power and cost information. The behavioural science literature on feedback-driven efficiency indicates that attributing explicit monetary costs to individual loads could encourage actions such as curtailing use of inefficient heaters, retiring high-draw legacy devices, and responding promptly to accidental-on states – interventions that also enhance electrical safety by flagging anomalous consumption patterns. Even a modest 5% reduction in the operating time of high-power appliances would yield appreciable financial savings and lower carbon emissions, far outweighing the monitor's own energy footprint of about 5–7 W (0.12–0.17 kWh day^{-1}, ≈ US$0.02). Subsequent design iterations could improve this net-positive balance even more by substituting ultra-low-power microcontrollers or implementing duty-cycled sensing, thereby reinforcing the system's sustainability benefits.

4.7. Limitations

While the proposed prototype effectively demonstrates the feasibility of edge-based appliance recognition, several limitations constrain its broader applicability. First, the current appliance library is limited in scope, and may not generalise well to devices with similar or overlapping load signatures. Second, the system operates under the assumption of non-overlapping appliance activity, lacking the capacity to disaggregate concurrent loads drawn from a single mains measurement. Third, the model does not yet incorporate mechanisms for online learning, which limits its adaptability to new devices or variations in electrical characteristics over time. In addition, the prototype functions purely as a passive monitoring tool and does not support any actuation or automated energy management. Last, the evaluation was conducted in controlled conditions without a long-term deployment or user-centred assessment, limiting insight into real-world usage patterns, interface effectiveness, and behavioural impact.

5. CONCLUSION

This study demonstrates an edge-based, AI-augmented energy-monitoring architecture that couples commercial hardware, Raspberry Pi, and non-intrusive current transformers with a custom-made neural-network classifier to furnish real-time, appliance-resolved power data via a web dashboard. Rigorous sensor calibration and careful data handling enabled 100% identification accuracy for four representative loads in controlled trials, thereby confirming that modest machine-learning models can achieve near-perfect performance when trained on discriminative features. By executing inference locally, the platform offers sub-second latency and preserves data privacy, validating key tenets of the non-intrusive load-monitoring and feedback literature that timely, granular information can catalyse behavioural efficiency gains. Although full-scale residential deployment would require expanded device libraries, multi-label separation, and online adaptation, the modular hardware and software stack readily accommodates additional sensors or more sophisticated algorithms. Consequently, the prototype is a realistic step towards cost-effective, consumer-oriented energy-management systems that are capable of evolving into intelligent assistants that not only diagnose consumption, but also orchestrate responsive control strategies to optimise cost and sustainability.

Future work will aim to improve its scalability and functionality. Priorities include expanding the appliance set, using deep features for better accuracy, enabling multi-load detection, adding online learning for adaptation, integrating smart control for automation, and running a long-term user study to assess impact and refine the interface.

ACKNOWLEDGEMENTS

The financial assistance of the National Research Foundation (NRF), the Automotive Industry Development Corporation (AIDC), and the Advanced Mechatronics Technology Centre (AMTC) at Nelson Mandela University with this research is hereby acknowledged. The opinions expressed and conclusions drawn are those of the authors, and do not necessarily reflect the views of the NRF, AIDC, or AMTC.

REFERENCES

[1] **D. Vine, L. Buys, and P. Morris**, "The effectiveness of energy feedback for conservation and peak demand: A literature review," Open Journal of Energy Efficiency, vol. 2, no. 1, pp. 7-15, 2013.

[2] **G.W. Hart**, "Nonintrusive appliance load monitoring," Proceedings of the IEEE, vol. 80, no. 12, pp. 1870-1891, 1992.

[3] **I. Antonopoulos et al.**, "Artificial intelligence and machine learning approaches to energy demand-side response: A systematic review," Renewable and Sustainable Energy Reviews, vol. 130, 109899, 2020.

[4] **A. Faustine, N.H. Mvungi, S. Kaijage, and M. Kisangiri**, "A survey on non-intrusive load monitoring methods and techniques for energy disaggregation problem," arXiv.1703.00785, 2017.

[5] **Z. Wang, W. Wang, Z. Zhang, F. Hu, X. Xia, and L. Chen**, "DeepEdge: A novel appliance identification edge platform for data gathering, capturing and labelling," Sensors, vol. 22, no. 7, 2432, 2022.

[6] **A. Zoha, A. Gluhak, M.A. Imran, and S. Rajasegarar**, "Non-intrusive load monitoring approaches for disaggregated energy sensing: A survey," Sensors, vol. 12, no. 12, pp. 16838-16866, 2012.

[7] **T.G. Hovgaard, L.F. Larsen, and J.B. Jørgensen**, "Robust economic MPC for a power management scenario with uncertainties," in 2011 50th IEEE Conference on Decision and Control and European Control Conference, 2011, pp. 1515-1520 .

[8] **N. Shabbir, K. Vassiljeva, H. Nourollahi Hokmabad, O. Husev, E. Petlenkov, and J. Belikov**, "Comparative analysis of machine learning techniques for non-intrusive load monitoring," Electronics, vol. 13, no. 8, 1420, 2024.

[9] **T. Ahmad and D. Zhang**, "Using the internet of things in smart energy systems and networks," Sustainable Cities and Society, vol. 68, 102783, 2021.

[10] **S. Khan, A.F. Latif, and S. Sohaib**, "Low-cost real-time non-intrusive appliance identification and controlling through machine learning algorithm," in 2018 International Symposium on Consumer Technologies (ISCT), 2018, pp. 32-36.

[11] **J. Winkler, H. Bousbiat, S. Jost, and W. Elmenreich**, "Energy disaggregation with NILM on a Raspberry Pi with smart-metering extension," in 2023 2nd International Conference on Power Systems and Electrical Technology (PSET), 2023, pp. 191-195.