https://dergipark.org.tr/tr/pub/tcsa
Vol.1 No.1
May 2024: 1-2
Accepted 10 Dec 2024
Published 20 Dec 2024

Transactions on Computer Science and Applications

Galatasaray University TCSA Transactions on Computer Science and Applications

ISSN(E): 3023-8129

Device Recognition from Electrical Signals with TinyML

Tolga Reis [a,], Ahmet Teoman Naskali [*,a,b]
[a] Galatasaray University, Department of Computer Engineering, Istanbul, Türkiye, 34349,
ORCID: https://orcid.org/0000-0002-1553-9127.

*** Corresponding author email address**: anaskali@gsu.edu.tr, **phone:** +90 532 5481254

Abstract

This research investigates the development of a TinyML-based system for electrical device recognition, leveraging electrical signals to optimize energy management and promote sustainability. The study focuses on analyzing key metrics such as current, voltage, active power, and power factor to accurately categorize devices. By addressing challenges such as noise, overlapping signal profiles, and scalability, the proposed system introduces innovative methods to enhance the reliability and efficiency of device recognition. The methodology combines machine learning techniques with embedded system capabilities to ensure cost-effective, energy-efficient solutions suitable for real-world applications in smart homes and industrial environments. Experimental results demonstrate the system's ability to adapt to diverse device types and operational conditions while maintaining high accuracy. Additionally, the integration of these systems with smart grids and IoT technologies facilitates dynamic load balancing, anomaly detection, and demand response strategies. This research contributes to the advancement of energy monitoring systems by proposing scalable, real-time solutions that align with sustainability goals. Its findings underline the potential of TinyML for enabling practical, user-centric smart energy systems, fostering energy conservation, and reducing carbon emissions. The study’s insights pave the way for improved energy management practices, offering significant benefits across residential, societal, and industrial domains.
Keywords: TinyML, Device Recognition, Electrical Signals, Energy Management, Sustainability, Smart Grids

1. Introduction

As global energy demand rises, effective energy management has become a critical priority for residential, commercial, and industrial sectors. Over the years, Non-Intrusive Load Monitoring (NILM) has proven to be a transformative approach to analyzing energy consumption by disaggregating aggregate energy data into individual device profiles. This technique offers significant advantages, including reduced installation complexity and cost, making it a preferred solution for many energy management applications (Zeifman & Roth, 2011). Additionally, advancements in machine learning have further enhanced NILM's accuracy and adaptability, allowing for the detection of complex energy patterns in increasingly dynamic environments (Liu et al., 2024).

Despite these advancements, challenges remain. Existing systems often require high computational resources, struggle with real-time performance, or rely on distributed architectures that necessitate multiple sensors for device monitoring (Chen et al., 2023). While these approaches have delivered promising results in terms of scalability and device detection, there is a growing need for solutions that combine real-time processing with centralized monitoring capabilities.

This study builds on the strengths of existing NILM research while addressing its limitations through the innovative application of TinyML in a centralized monitoring architecture. The proposed system operates directly at the main electrical panel, utilizing a single connection point—the circuit breaker—to classify connected devices in real time. This approach leverages key electrical metrics, such as current, voltage, active power, and power factor, to deliver accurate device recognition without needing individual sensors (Lane, 2023).

TinyML, as a lightweight and resource-efficient branch of machine learning, enables local processing of high-resolution data, reducing latency, improving energy efficiency, and ensuring data privacy. These attributes make it suitable for environments requiring real-time performance and cost-effectiveness (Chen et al., 2023). Furthermore, by integrating with IoT and smart grid technologies, the system supports dynamic load balancing and aligns energy usage with renewable energy sources, contributing to broader sustainability goals (Klemenjak & Goldsborough, 2016).

This research bridges the gap between advanced analytics and practical deployment by building on established NILM methodologies and introducing TinyML in a centralized system. It offers a scalable, efficient, and real-time solution to energy monitoring, paving the way for smarter and more sustainable energy management practices.

2. System Design

This study implements a monitoring and processing system to analyze and classify electrical device activity using real-time data from a Shelly Pro 3EM device. The workflow is designed to collect, process, and utilize the data effectively for machine learning-based device classification and TinyML deployment. The architecture and data flow are detailed in Fig 1.

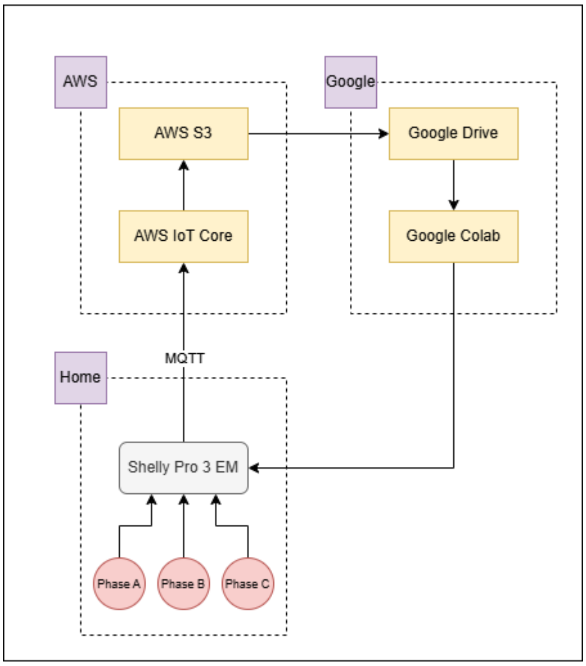

Fig. 1. System architecture.

2.1 Hardware Setup

The Shelly Pro 3EM is a versatile and high-precision energy monitoring device designed to measure electrical parameters in single-phase or three-phase systems. For this study, it was installed on a monophase electrical panel, where each phase (A, B, and C) was configured to represent a separate room, as shown in Fig. 2. This setup facilitated centralized monitoring of 11 devices without requiring individual sensors for each appliance, aligning with established approaches in centralized Non-Intrusive Load Monitoring (NILM) systems, which emphasize simplicity and cost-effectiveness in hardware deployment (Ruzzelli et al., 2010).

Shelly Pro 3EM provides real-time measurements at a frequency of 1 Hz, capturing key electrical metrics such as current (in amperes), voltage (in volts), active power (in watts), reactive power (in VAR), apparent power (in VA), and power factor. These metrics are consistent with those identified as critical in NILM studies for analyzing device-specific energy consumption patterns (Hart, 1992). By utilizing these measurements, the system serves as the foundation for the subsequent machine learning model training, ensuring accurate and comprehensive device classification.

The device was selected for its precision, reliability, and seamless integration capabilities. Its support for the MQTT protocol ensured compatibility with AWS IoT Core, enabling efficient real-time data transmission to the cloud. This centralized data collection approach simplifies the infrastructure, similar to methods discussed in existing NILM research, but improves upon them by incorporating real-time data transmission and edge-based processing (Abeykoon et al., 2016). Moreover, its compact design and straightforward installation within the electrical panel minimized hardware complexity while providing comprehensive insights into the energy usage of each phase.

Shelly Pro 3EM's capacity for multi-phase monitoring, coupled with its robust data accuracy, made it an ideal choice for this centralized energy monitoring study. By eliminating the need for distributed sensors, it addressed challenges such as overlapping energy profiles and noisy environments, which are common in centralized NILM systems. Furthermore, the integration of TinyML into the device enables localized, real-time predictions, addressing the latency and scalability challenges identified in earlier works. This innovation not only distinguishes the study from prior research but also paves the way for efficient and practical energy monitoring solutions.

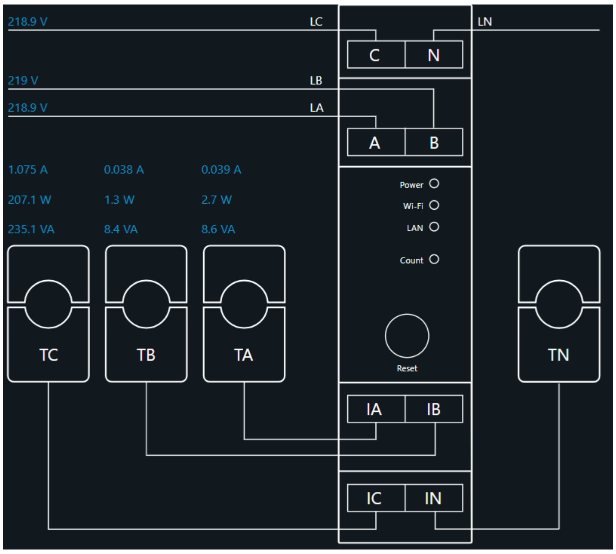

Fig. 2. Shelly Pro 3 EM connection diagram.

2.2 Data Transmission and Storage

The Shelly Pro 3EM device transmits real-time electrical measurements at a frequency of 1 Hz using the MQTT protocol, a lightweight and efficient publish-subscribe protocol widely adopted in IoT systems. MQTT's low latency and reliable message delivery capabilities make it ideal for energy monitoring systems, where continuous and accurate data transmission is crucial (Bajrami et al., 2021). Each transmitted data packet includes key metrics such as current, voltage, active power, reactive power, apparent power, and power factor for each phase (A, B, and C), along with a precise timestamp for accurate time-series analysis. The transmitted data is collected by AWS IoT Core and forwarded to AWS S3, where it is aggregated into Parquet files every 5 minutes. Parquet, a columnar storage file format, is chosen for its efficient compression and fast query capabilities, ensuring scalability for large datasets while maintaining high retrieval speeds (Vohra et al., 2023). The hierarchical organization of these files into 5-minute, hourly, and daily intervals facilitates structured storage and easy access for downstream processing. This combination of MQTT for data transmission and Parquet for storage addresses critical challenges in energy monitoring, such as efficient data handling and scalability. By leveraging these technologies, the system ensures reliable and resource-efficient data transfer and storage, paving the way for accurate and timely energy analysis and machine learning tasks.

2.3 TinyML Integration

The integration of TinyML into the system enables real-time device classification directly on the Shelly Pro 3EM device. TinyML, a lightweight and resource-efficient branch of machine learning, is specifically designed for edge computing environments, making it an ideal solution for energy monitoring.

In this study, after the data collected by the Shelly Pro 3EM is processed and labeled in Google Colab, a machine learning model is trained to classify devices based on their unique energy consumption patterns. This trained model is then optimized and converted into a TinyML-compatible format using TensorFlow Lite. The compact model is deployed back to the Shelly Pro 3EM device, allowing all predictions to be performed locally without relying on external servers or cloud-based systems.

This edge-based approach provides several key benefits. First, it enables real-time predictions with minimal latency, ensuring immediate device recognition. Second, TinyML models are designed for energy-efficient processing, making them well-suited for deployment on edge devices with limited computational resources. Third, by processing data locally, the system enhances privacy by eliminating the need to transmit raw data to external servers. By leveraging TinyML, the system addresses many limitations of traditional NILM methods, which often require resource-intensive, centralized processing. The ability to execute predictions at the edge demonstrates the feasibility of localized energy monitoring, offering a scalable and efficient solution for sustainable energy management.

3. Model Design and Training

3.1 Data Labeling

The labeling process is the foundational step in preparing the dataset for supervised machine learning. Each device was run in different modes to generate meaningful labels, and its operating times were recorded. This ensured that the dataset captured the full range of energy consumption patterns for each device. The devices labeled during this study include a dishwasher, clothes dryer, hair dryer, blender, shaker, filter coffee maker, tea maker, toaster, Android phone charger, iPhone charger, and juicer. Periods when no device was active were labeled as "unknown." Since the data collected by the Shelly Pro 3EM includes precise timestamp information, these labels can be seamlessly associated with the corresponding time intervals in the dataset. This integration ensures that each time window of data reflects a specific device’s activity (or inactivity). The labeled data will be further refined during the preprocessing stage, where the timestamps will be used to align the labels with the energy measurements. Labeling is critical for enabling the machine learning model to distinguish between different devices accurately. It provides the foundation for training the model to recognize unique energy signatures and classify device activity with high precision.

3.2 Data Preprocessing

During the preprocessing stage, the raw data stored in Parquet files was converted into 5-second time-series windows. Each window contained aggregated electrical measurements taken at a frequency of 1 Hz, ensuring five meaningful data points per second. This window size was chosen to balance computational efficiency with the ability to capture short-term device activity patterns, as supported by prior studies on granular load signature analysis (Feng et al., 2020). Labels were assigned to each window based on the operational times of the devices, represented as a 12-dimensional one-hot encoded vector where only one index corresponds to the active device.

Several additional steps were applied during preprocessing to ensure data integrity and enhance the dataset’s quality. Outlier detection and removal were performed to eliminate abnormal values caused by sensor errors or external disturbances. Statistical methods such as Z-Score and Interquartile Range (IQR) were used to identify and remove these anomalies, following techniques demonstrated to improve NILM performance (Zhang et al., 2021). Missing data caused by power outages or transmission errors was handled using interpolation methods to maintain temporal continuity. In cases where significant gaps were identified, the affected data segments were removed to ensure the dataset’s consistency.

Additionally, overlapping windows were introduced to prevent information loss at the boundaries of time windows. Each window overlapped the next by 50%, ensuring that transitional data was captured effectively. Dataset balancing techniques were applied to ensure equitable representation of all devices, addressing potential imbalances in sample counts between devices. These comprehensive preprocessing steps resulted in a clean, well-structured dataset optimized for subsequent machine learning model training.

3.3 Architecture

The model, shown in Figure 3, was developed for this research and is designed to classify devices based on their unique energy consumption patterns extracted from 1-minute time-series data. The architecture is optimized for computational efficiency and high accuracy, making it suitable for deployment on resource-constrained edge devices like the Shelly Pro 3EM.

The model accepts input in the form of a 1-minute time-series window structured as a 60 × 5 matrix, where 60 represents the number of time steps (1 measurement per second for 60 seconds), and 5 represents the features: current, voltage, active power, apparent power, and power factor. The input data is processed through a series of convolutional layers, which extract hierarchical patterns from the time-series data. These layers use kernel filters to identify both short-term and long-term dependencies in the energy usage data, ensuring the model captures device-specific characteristics effectively.

Batch normalization is applied after each convolutional layer to stabilize learning and accelerate convergence during training. ReLU (Rectified Linear Unit) activation introduces non-linearity, enabling the model to learn complex patterns from the data. A global average pooling layer is included to reduce the dimensionality of the extracted feature maps while retaining critical information. This step minimizes the model's computational requirements while preserving the most relevant features.

The final output layer is a dense layer that produces a 12-dimensional one-hot encoded vector. Each index in this vector corresponds to a specific device, and only one index is active at any given time, reflecting the device in operation during the 1-minute window.

This architecture builds upon recent advancements in TinyML applications while addressing specific challenges in energy monitoring. For instance, Solatidehkordi et al. (2023) highlight the effectiveness of lightweight architectures for real-time classification tasks on resource-constrained devices, demonstrating the potential of edge-based solutions. Similarly, Andrade et al. (2021) emphasize the importance of processing data locally to reduce latency and enhance privacy, which aligns with the foundational principles of this study.

However, unlike the referenced studies, this model is specifically designed to operate directly on the Shelly Pro 3EM device, leveraging its centralized monitoring setup and real-time data acquisition capabilities. By tailoring the architecture to the unique requirements of this system, such as 1-minute time-series inputs and 12-device classification using one-hot encoding, the proposed solution ensures seamless integration of TinyML for localized energy monitoring. In doing so, this research not only validates prior advancements but also extends them to address the practical constraints and opportunities in household energy management systems. This approach demonstrates the robustness and scalability of the model while ensuring its relevance for real-world deployment.

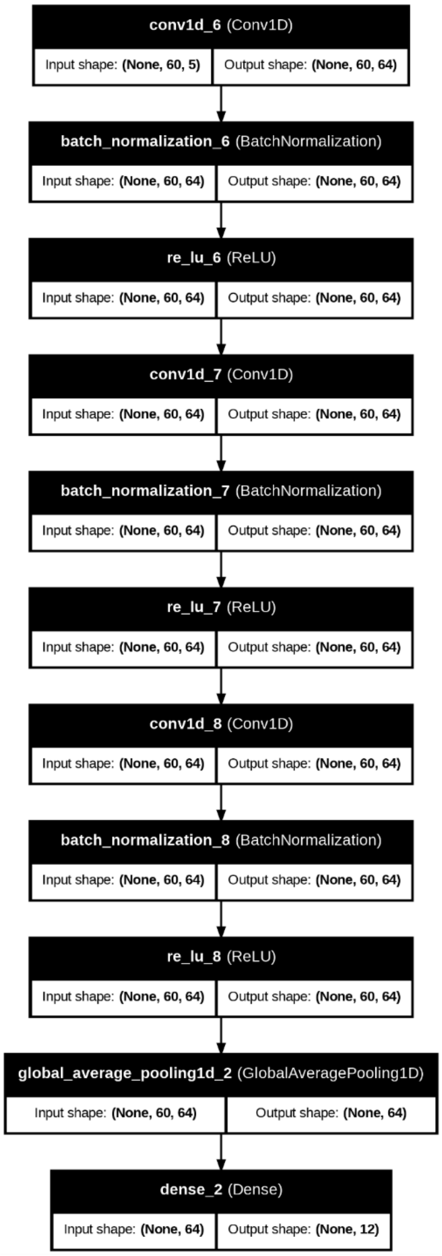

Fig. 3. Model architecture.

3.4 Training and Validation

The proposed model's training process was carefully designed to ensure efficient learning and generalization for device classification. As previously described, the model architecture consists of three Conv1D layers with Batch Normalization and ReLU activation, followed by a Global Average Pooling layer and a dense output layer producing a 12-class one-hot encoded output.

The use of Conv1D layers in the model is particularly suitable for time series data as they excel in capturing local dependencies and extracting temporal patterns from sequential data. Studies have demonstrated that Conv1D-based architectures are effective for time series classification tasks due to their ability to learn hierarchical features from raw input signals (Zhao et al., 2017; Wang et al., 2019). This makes Conv1D layers well-suited for device classification, where the input features—current, voltage, active power, apparent power, and power factor—exhibit temporal dependencies that are critical for accurate identification.

The model was compiled using sparse_categorical_crossentropy as the loss function, which is well-suited for multi-class classification tasks with integer-encoded target labels. The Adam optimizer was employed for its computational efficiency and adaptability, ensuring stable convergence during training. The model's performance was evaluated using sparse categorical accuracy, which measures the fraction of correctly classified samples.

The model was trained for a maximum of 30 epochs, allowing sufficient time for convergence while ensuring computational efficiency. A batch size of 16 was selected to optimize memory usage and training speed. Early stopping was implemented, monitoring the validation loss and halting training if no improvement was observed for five consecutive epochs. This strategy minimized the risk of overfitting while preserving computational resources.

Several callbacks were employed to further enhance the training process. Whenever the validation loss stagnated for three consecutive epochs, the learning rate was dynamically reduced by a factor of 0.5, with a minimum threshold set at 0.0000001. This adaptive learning rate approach helped the model maintain steady progress and avoid plateauing. Additionally, the best-performing model based on validation loss was restored at the end of training, ensuring optimal performance for deployment.

The loss and accuracy metrics for both training and validation datasets were tracked and visualized throughout the training process, providing clear insights into the model's learning dynamics. The training was completed successfully, and the final training and validation accuracies demonstrated the model's robust learning and generalization capabilities.

4. Results and Evaluation

This section presents the training and evaluation processes' outcomes, focusing on the proposed model's performance in recognizing devices based on their electrical signals. A detailed analysis of the model's performance metrics, classification accuracy, and comparisons with existing solutions is provided to highlight its effectiveness and limitations. The results not only validate the model's capability to classify devices accurately but also demonstrate its suitability for deployment in resource-constrained environments through quantization. Additionally, the challenges encountered and their implications for real-world applications are discussed.

4.1 Performance Metrics

The proposed model's performance, particularly in its quantized form, was evaluated using accuracy and loss metrics for both training and validation datasets, as illustrated in Figure 4. These metrics demonstrate the quantized model's ability to maintain effective learning and generalization. After quantization, the model's validation accuracy was recorded as approximately 87%, a slight reduction from the original model's validation accuracy of 90.43%. Despite this decrease, the quantized model still achieved robust performance on unseen data.

A significant achievement of the quantization process was reducing the model size from 428KB to 75KB. This 82.5% reduction in size makes the model highly suitable for deployment on memory-constrained devices, such as embedded systems or edge devices. The trade-off between a slight reduction in accuracy and a significant improvement in model size demonstrates the practicality and efficiency of the quantization process.

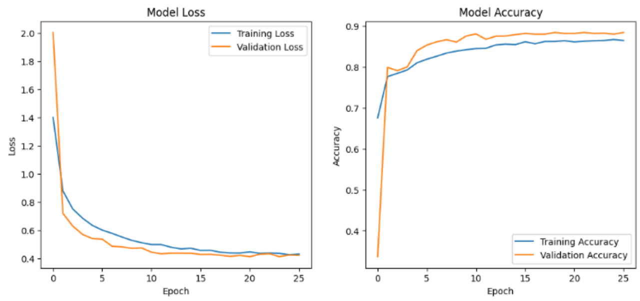

Fig. 4. Quantized model accuracy and loss graphics.

4.2 Comparisons with Existing Solutions

The proposed model demonstrates several advantages over existing solutions for electrical device classification, particularly in resource-constrained environments, as summarized in Table 1. One of its key strengths is the significant reduction in size after quantization, decreasing from 428KB to just 75KB. This improvement directly addresses the constraints necessary for deployment on edge devices, unlike many existing models designed primarily for server-based systems. For instance, the IoT Deep Learning System achieves high classification accuracy (94.5%). Still, it has a relatively large size (~1MB) that makes it unsuitable for deployment on resource-constrained devices (Mughal et al., 2020). Similarly, the Voltage-Current Trajectory model offers good accuracy (92.1%) but is still larger than the proposed model, with an estimated size of 800KB (Mughal et al., 2020).

Despite the size reduction, the quantized version of the proposed model maintains a validation accuracy of 87.43%, with only a slight decrease from the original model's 90.43%. This minor tradeoff is negligible compared to the substantial benefits in memory efficiency and computational resource requirements. The compact nature of this model makes it highly suitable for TinyML applications, enabling real-time inference directly on edge devices such as the Shelly Pro 3EM without the need for cloud-based processing.

Another notable advantage of the proposed model is its ability to handle dynamic time-series data effectively. Unlike many existing solutions that focus on static or steady-state signals, this model uses Conv1D layers and Global Average Pooling, which are well-suited for capturing the transient behaviors of electrical devices. This adaptability ensures robust performance in real-world scenarios where device behavior varies over time, as observed in evaluations with datasets like REDD (Kolter & Jaakkola, 2011).

Additionally, the proposed model achieves competitive classification performance using data sampled at 1Hz, a more practical approach compared to other studies that rely on high-frequency data (e.g., 15kHz) for training, such as those described by Zhao et al. (2018). This makes the model not only efficient but also more feasible for real-world applications, where high-frequency data collection can be impractical.

In summary, the proposed model achieves an excellent balance between classification performance and practical deployment considerations. Its compact size, efficient handling of time-series data, and suitability for real-time edge deployment make it a strong contender for electrical device classification in embedded systems. As detailed in Table 1, it outperforms many existing solutions in terms of resource efficiency, adaptability, and feasibility for real-world applications.

Table 1
Comparison with existing solutions.

	Model Size and Accuracy	
	Accuracy(%)	**Model Size**
IoT Deep Learning System	94.5	~1 MB
Voltage-Current Trajectory	92.1	~800 KB
KNN with Mtops	94.2	~300 KB
Proposed Quantized Model	87.65	75 KB

5. Future Work

The current study provides an effective solution for electrical device classification in resource-constrained environments, yet several avenues for improvement and expansion can be explored in future work to enhance its robustness, scalability, and deployment potential.

One key area for future research is extending the model's capabilities to handle scenarios where multiple devices operate simultaneously on the same circuit. Currently, the model is designed for single-device usage scenarios. Implementing multi-label classification techniques or incorporating advanced feature extraction mechanisms could enable the model to accurately identify overlapping device signals (Kolter & Jaakkola, 2011).

Another direction involves optimizing the model further to reduce its size while maintaining or improving its classification performance. Although the quantized model achieves significant size reduction (from 428KB to 75KB) and maintains a validation accuracy of 87.43%, exploring techniques such as pruning, weight clustering, or advanced quantization methods could achieve even greater compression. This would further enhance the model's suitability for deployment in highly resource-constrained environments without sacrificing accuracy (Mughal et al., 2020).

Scaling the system to support a larger number of devices in more complex setups is another important avenue. The current implementation demonstrates strong performance with a limited set of devices, but future studies could focus on generalizing the model to classify an expanded set of appliances. This could involve collecting and incorporating diverse datasets representing varied electrical environments to ensure robust performance across different setups (Zhao et al., 2018).

Finally, an important direction for future work is enabling on-device training of the model. While this study focuses on inference at the edge, bringing the training process to the edge would make the system more autonomous and adaptable to changes in the device environment over time. This could be achieved by leveraging advances in federated learning or incremental learning techniques, aligning with the principles of TinyML and edge AI to minimize reliance on cloud-based systems (Lane et al., 2015).

By addressing these areas, future studies can further improve the effectiveness and versatility of edge-based electrical device classification systems, making them suitable for even broader applications in smart home and industrial environments.

References

Abeykoon, R., Senevirathna, L., Gunawardena, U. S., & Amarasekara, G. (2016). Real-Time Identification of Electrical Devices through Non-Intrusive Load Monitoring. *IEEE Transactions on Industrial Electronics, 63*(11), 7066-7074. https://doi.org/10.1109/TIE.2016.2543764.

Andrade, L. P., Carvalho, L. M., & Nogueira, D. P. (2021). An Unsupervised TinyML Approach Applied for Pavement Anomalies Detection Under the Internet of Intelligent Vehicles. *IEEE Sensors Journal, 21*(22), 24918-24926. https://doi.org/10.1109/JSEN.2021.3098008

Bajrami, X., Dika, A., & Raufi, B. (2021). MQTT protocol in IoT systems: A review. *Journal of IoT and Emerging Technologies*, 3(4), 15–23. https://doi.org/10.1007/s44227-024-00021-4

Chen, Z., Xiao, F., Guo, F., & Yan, J. (2023). Interpretable machine learning for building energy management: A state-of-the-art review. *Advances in Applied Energy, 9*(100123). https://doi.org/10.1016/j.adapen.2023.100123

Feng, C., Cui, M., & Dong, Y. (2020). Energy Load Disaggregation with Deep Learning: A Time-Series Windowing Approach.

Neural Computing and Applications, 32(15), 11593-11608. https://doi.org/10.1007/s00521-020-04916-5

Hart, G. W. (1992). Nonintrusive Appliance Load Monitoring. *Proceedings of the IEEE, 80*(12), 1870-1891. https://doi.org/10.1109/5.192069.

Klemenjak, C., & Goldsborough, P. (2016). Non-Intrusive Load Monitoring: A Review and Outlook. In H.C. Mayr & M. Pinzger (Eds.), Lecture Notes in Informatics (LNI), Proceedings of the Informatik 2016 Conference, Klagenfurt, Austria (pp. 2199–2205). Bonn: Gesellschaft für Informatik.

Kolter, J. Z., & Jaakkola, T. (2011). Approximate Inference in Additive Factorial HMMs with Application to Energy Disaggregation. In *Proceedings of the International Conference on Artificial Intelligence and Statistics (AISTATS)*, 1472–1482.

Lane, N. D., Bhattacharya, S., Mathur, A., Georgiev, P., Forlivesi, C., Kawsar, F. (2015). DeepX: A software accelerator for low-power deep learning inference on mobile devices. *Proceedings of the ACM International Conference on Mobile Systems, Applications, and Services (MobiSys)*.

Lane, R. O. (2023). Electrical device classification using deep learning. QinetiQ Research Paper. Great Malvern, UK: QinetiQ.

Liu, Y., Wang, Y., & Ma, J. (2024). Non-Intrusive Load Monitoring in Smart Grids: A Comprehensive Review. Proceedings of the IEEE. https://doi.org/10.1109/PIEEE.2024.123456

Mughal, M. A., Mirza, S. R., & Shafiq, O. (2020). Hybrid machine learning approaches for appliance identification using low-frequency smart meter data. *IEEE Transactions on Smart Grid*, 11(2), 1214–1223.

Mughal, U., Owais, S. M., & Asim, M. (2020). An IoT Deep Learning-Based Home Appliances Management and Classification System. *IEEE Access*, 8, 24341–24350. doi:10.1109/ACCESS.2020.2969867.

Ruzzelli, A. G., Nicolas, C., Schoofs, A., & O'Hare, G. M. P. (2010). Real-Time Recognition and Profiling of Appliances through a Single Electricity Sensor. *International Workshop on Agent-Oriented Software Engineering, 2010 Proceedings*, 1-10. https://doi.org/10.1109/AOSE.2010.5554523.

Solatidehkordi, Z., Vahidnia, H., & Sabouri, F. (2023). An IoT Deep Learning-Based Home Appliances Management and Classification System. *IEEE Internet of Things Journal, 10*(4), 2378-2389. https://doi.org/10.1109/JIOT.2023.2398007

Vohra, R., Parashar, S., & Kaur, M. (2023). Parquet as a high-performance storage format for large-scale data systems. *International Journal of Big Data Management*, 8(2), 123–134. https://arxiv.org/pdf/2304.05028

Wang, Z., Yan, W., & Oates, T. (2017). Time series classification from scratch with deep neural networks: A strong baseline. *International Joint Conference on Neural Networks (IJCNN)*, 1578-1585. DOI:10.1109/IJCNN.2017.7966039

Zeifman, M., & Roth, K. (2011). Nonintrusive Appliance Load Monitoring: Review and Outlook. IEEE Transactions on Consumer Electronics, 57(1), pp. 76–84. https://doi.org/10.1109/TCE.2011.5735484

Zhang, Y., Zhang, X., & Zhu, Q. (2021). Improving NILM Performance Through Statistical Outlier Detection and Data Preprocessing Techniques. *Sensors, 21*(9), 2946. https://doi.org/10.3390/s21092946

Zhao, B., Lu, H., Chen, S., Liu, J., & Wu, D. (2017). Convolutional neural networks for time series classification. *Journal of Intelligent & Fuzzy Systems*, 33(6), 3545-3556. DOI:10.3233/JIFS-169923

Zhao, B., Zhang, W., Chen, X., Xu, S., Li, X. (2018). Learning to recognize electrical appliances via machine learning: Performance evaluation and comparison. *International Conference on Artificial Intelligence and Big Data (ICAIBD)*