Journal of Urban Development and Smart Cities
2025, 2(1), 142-150

LOW-FREQUENCY NON-INTRUSIVE LOAD MONITORING FOR RESIDENTIAL ENERGY MANAGEMENT IN SMART CITIES USING A MODIFIED K-NEAREST NEIGHBOUR ALGORITHM

Adson Damasceno
Ivanovitch Silva

Residential and small commercial buildings remain central to urban electricity demand, yet appliance-level visibility is often limited by the cost and installation burden of intrusive sub-metering. This study presents a machine-learning framework for non-intrusive load monitoring (NILM) that recovers appliance-level consumption patterns from a single smart-meter stream using steady-state active power. The framework combines a Bagging regressor for load prediction with a modified multiclass K-Nearest Neighbour (KNN) classifier for appliance identification and is evaluated on the Dutch Residential Energy Dataset (DRED), a low-frequency 1 Hz dataset containing aggregate and appliance-level measurements from a monitored household in the Netherlands. Using a 70% training split and 30% testing split, zero-imputation for sparse missing values, and time-derived features capturing hourly and daily usage variation, the proposed pipeline achieves strong predictive performance while preserving computational efficiency suitable for smart-meter deployment. The Bagging regressor records the best regression performance among the evaluated models, with a mean squared error of 99.3152, root mean squared error of 9.9657, mean absolute error of 0.2813, and coefficient of determination of 0.9624. For appliance classification, the optimised multiclass KNN model achieves a precision of 0.7497, recall of 0.7825, F_1-score of 0.7531, and overall accuracy of 78.25%. Descriptive analysis of the monitored period further shows consistent refrigerator demand, peak aggregate usage on 6 and 13 July, and the lowest household demand between midnight and 05:33. These results support the use of low-frequency NILM as a practical and scalable component of urban energy analytics, with direct relevance for smart-home feedback, demand-side management, infrastructure planning, and carbon-aware city operations.

Index Terms — non-intrusive load monitoring; smart cities; residential energy analytics; demand-side management; smart meters; K-nearest neighbour; bagging regressor

https://doi.org/10.66033/judsc2025-215

INTRODUCTION

Urban development increasingly depends on intelligent energy systems that can support demand-side management, reduce waste, and improve planning across residential and small commercial buildings. A persistent challenge is that most buildings are metered only at the aggregate level, which obscures the contribution of individual appliances to total demand. This lack of granularity limits consumer feedback, slows fault detection, and constrains the quality of operational data available for distribution planning and energy-efficiency policy.

Non-intrusive load monitoring (NILM), also known as energy disaggregation, addresses this problem by inferring appliance-level demand from a single measurement point rather than requiring a sensor on every device. In contrast to intrusive load monitoring, NILM offers a lower-cost architecture that is especially attractive for large-scale deployment in urban settings where hardware additions must remain minimal. Classical NILM research established the feasibility of this approach through appliance signature analysis and event-based disaggregation [1]. More recent work has explored machine-learning and deep-learning formulations, but many of these approaches require high-frequency data, substantial computation, or extensive model complexity.

The present study develops a practical NILM framework tailored to the operating constraints of conventional smart meters. The central design principle is to use only steady-state active power sampled at low frequency, thereby preserving compatibility with installed metering infrastructure while still enabling both regression-based load prediction and multiclass appliance identification. The method combines a Bagging regressor with a modified K-Nearest Neighbour (KNN) classifier and evaluates performance on the Dutch Residential Energy Dataset (DRED).

This framing aligns directly with the needs of a smart-cities and urban development readership. Accurate low-cost energy disaggregation can support better household energy behaviour, identify inefficient or degraded appliances, enhance utility-side planning for expansions and renovations, and contribute to long-term reductions in system-wide carbon emissions. In this sense, NILM is not merely a device-level analytics problem; it is an urban infrastructure intelligence problem.

LITERATURE CONTEXT

NILM research has historically relied on two broad methodological families: electrical-signature engineering and model-based learning. Early studies established that appliance identity can be inferred from changes in voltage, current, active power, and reactive power measured at a single service entry point [1]. Subsequent work extended these ideas using probabilistic models, decision trees, nearest-neighbour methods, and, more recently, deep neural networks.

A recurring trade-off in the literature concerns *accuracy versus deployability*. High-frequency transient approaches can improve discrimination across devices with similar steady-state demand, but they also require faster sampling, more computation, and more complex feature extraction. Deep-learning models often offer strong benchmark performance, yet their computational footprint can limit integration into edge devices and low-cost smart-meter environments.

Low-frequency NILM therefore remains highly relevant to urban-scale implementation. In this regime, KNN is attractive because of its simplicity, non-parametric flexibility, and ability to accommodate non-linear class boundaries when an appropriate distance metric is used [2]. Ensemble regression methods, including Bagging, are similarly attractive because they reduce overfitting, improve generalisation, and remain computationally feasible relative to heavier architectures [3]. The key research task is thus to obtain strong predictive performance while retaining the practical advantages required for city-scale deployment through existing metering infrastructure.

MATERIALS AND METHODS

System objective and application setting

The study targets appliance-level monitoring in residential and small commercial settings using a single smart meter. The system is designed to: (i) predict aggregate and appliance-related consumption behaviour, (ii) classify connected appliances from steady-state active power signatures, and (iii) provide information that can support energy cost control, appliance fault awareness, and urban demand-side planning.

Dataset

Model development is based on the Dutch Residential Energy Dataset (DRED), an open-source dataset collected from a household in the Netherlands. DRED contains both aggregated and appliance-specific active power measurements acquired continuously over a six-month period from 5 July to 5 December 2015 at a sampling frequency of 1 Hz. The acquisition setup includes house-level and appliance-level measurements, and the dataset also provides ambient, occupancy, and household information in separate CSV files.

Table 1: Dataset and preprocessing summary.

Item	Description
Primary dataset	Dutch Residential Energy Dataset (DRED)
Monitoring horizon	5 July 2015 to 5 December 2015 (approximately six months)
Sampling frequency	1 Hz
House-level acquisition	Smart meter and Plugwise Smile; active power and time
Appliance-level acquisition	Plugwise Circle; active power and time
Auxiliary information	Ambient, occupancy, and household CSV files
Data availability	Above 90%–95% for aggregate and appliance data
Dropout rate	Less than 5% for active power data
Training/testing split	70% training, 30% testing
Missing-data treatment	Zero imputation
Engineered temporal features	Hourly and daily variations derived from steady-state active power

The dataset was selected because of three practical advantages: high availability, low dropout, and stable baseline monitoring of occasionally used appliances. These features make it a suitable benchmark for evaluating low-frequency NILM in real building contexts.

Preprocessing

Because active-power continuity is critical for both disaggregation and forecasting, missing observations were handled through zero imputation. Given the dataset's reported dropout rate of less than 5%, this choice preserves the observed structure of household consumption without materially distorting the operating state of appliances. The preprocessing pipeline also extracts time-based features from steady-state active power, specifically hourly and daily variation patterns, to improve model discrimination and better capture household usage regularity.

Load categories

The monitored appliances span the major operational categories commonly used in NILM:

- Type I (on/off): appliances with binary states, such as television, fan, oven, microwave, toaster, and sockets.
- Type II (finite-state): appliances with multiple discrete operating states, including the washing machine and electric heating element.
- Type III (continuously variable): appliances whose power can vary with user control, such as the cooker and electric heating element.
- Type IV (permanent): appliances that draw continuous power, including the fridge, television, and laptop computer.

This operational diversity makes DRED a meaningful testbed for low-frequency appliance classification.

Modified KNN classifier

The classification component uses a modified multiclass KNN model tailored to appliance recognition from steady-state active power. KNN assigns a class label to an unknown observation by examining the closest labelled instances under a chosen distance metric. The Euclidean distance used in the final tuned model is given by

$$d(\mathbf{x},\mathbf{z}) = \sqrt{\sum_{p=1}^{P} (x_p - z_p)^2}, \tag{1}$$

where P is the feature dimension and $\mathbf{x}$ and $\mathbf{z}$ are feature vectors.

To express class membership likelihood, the model uses inverse-distance information so that observations closer to the query point exert stronger influence. In formal terms, the inverse-distance contribution can be written as

$$u(\mathbf{x},\mathbf{z}) = \frac{1}{d(\mathbf{x},\mathbf{z})}. \tag{2}$$

The overall optimisation process considered the number of neighbours, weighting scheme, power parameter, distance metric, algorithm, and leaf size. GridSearchCV was used to identify the best-performing configuration.

Table 2: Selected hyperparameters for the optimised multiclass KNN classifier.

Hyperparameter	Selected value
Number of neighbours	21
Weighting scheme	Uniform
Power parameter	2
Distance metric	Euclidean
Algorithm	Auto
GridSearchCV folds	20

The final model therefore preserves KNN's computational simplicity while adapting its operating parameters to the empirical structure of the DRED data.

Bagging regressor

The predictive component uses a Bagging regressor to estimate power consumption patterns. Bagging aggregates multiple bootstrapped learners to improve stability and reduce variance, making it a strong candidate for household load prediction where consumption behaviour is non-linear and heterogeneous. For comparison, the study also evaluates KNN regression, Decision Tree regression, Extreme Gradient Boosting, and Linear Regression.

Regression performance is assessed using mean squared error (MSE), root mean squared error (RMSE), mean absolute error (MAE), and the coefficient of determination (R^2).

RESULTS

Descriptive load patterns

The descriptive analysis focused on a highlighted period from 5 to 17 July 2015 within the broader DRED monitoring horizon. Aggregate power plots showed the highest daily peaks on 6 July and 13 July, while the lowest daily aggregate consumption occurred on 11 July. The hourly aggregate profile indicated the lowest load conditions between 00:00 and 05:33:20, followed by a steady rise after that interval and a pronounced peak at approximately 13:00.

Appliance-level inspection showed that the fridge was the most consistently active device throughout the observed period, with a stable daily consumption profile and small fluctuations attributable to defrost cycles. The cooker showed major activity on days 6 and 13, while peak usage for sockets, the washing machine, and the toaster occurred on days 9, 15, and 16, respectively. Taken together, these patterns indicate that the refrigerator imposes the largest persistent household cost burden, whereas several other appliances contribute shorter-duration but sharper demand spikes.

Table 3: Observed consumption patterns in the monitored household.

Observed feature	Interpretation
Highest daily aggregate peaks	6 July and 13 July
Lowest daily aggregate consumption	11 July
Lowest hourly demand interval	00:00 to 05:33:20
Highest hourly aggregate peak	Approximately 13:00
Most consistently active appliance	Fridge
Fridge profile	Continuous use with minor fluctuations linked to defrost cycles
Cooker peak days	6 July and 13 July
Sockets peak day	9 July
Washing machine peak day	15 July
Toaster peak day	16 July

Regression performance

Among the evaluated regression models, the Bagging regressor produced the strongest overall performance. Its R^2 score of 0.9624 indicates that it explains the vast majority of the variance in the target signal, while its RMSE of 9.9657 is the lowest among all candidate models.

Table 4: Performance of selected regression models.

Regression model	MSE	RMSE	MAE	$\mathbf{R^2}$
K-Nearest Neighbour	402.7246	20.0680	2.3089	0.7533
Decision Tree	132.5814	11.5144	0.2211	0.9470
Extreme Gradient Boosting	123.6677	11.1206	1.3843	0.9170
Bagging	99.3152	9.9657	0.2813	0.9624
Linear Regression	2020.1429	44.9460	13.2874	0.1043

This outcome is important in the context of urban energy analytics. The Bagging model offers a strong predictive signal without the cost, data requirements, or computational burden commonly associated with high-complexity approaches. Its ensemble structure also reduces overfitting and improves generalisation across varying appliance-use patterns.

Baseline appliance classification performance

For appliance identification, the study compared five classification models. The untuned KNN classifier outperformed the other evaluated baselines, recording the highest accuracy and recall among the tested models.

Table 5: Performance of selected classification models.

Classification model	Precision	Recall	$\mathbf{F_1}$ score	Accuracy
K-Nearest Neighbour	0.6991	0.7751	0.6845	0.7751
Logistic	0.5965	0.7724	0.6732	0.7724
AdaBoost	0.5965	0.7721	0.6730	0.7721
Bagging	0.6052	0.7296	0.6596	0.7296
Decision Tree	0.6045	0.5751	0.5894	0.5751

The classification boundary was reported to be non-linear, which further supports the suitability of KNN for this task. Unlike linear methods, KNN can better accommodate irregular class separation when appliances exhibit overlapping or curved load patterns in feature space.

Optimised multiclass KNN performance

After hyperparameter tuning, the optimised KNN classifier improved its accuracy by approximately 1.1 percentage points over the untuned baseline. The final model correctly classified appliances 78.25% of the time, with corresponding gains in precision and F_1-score.

Table 6: Performance of the optimised multiclass KNN model.

Model	Precision	Recall	$\mathbf{F_1}$ score	Accuracy
Optimised K-Nearest Neighbour	0.7497	0.7825	0.7531	0.7825

Although approximately 21.75% of instances remained misclassified, the overall result is strong given that the model relies only on steady-state active power and must separate appliances with overlapping signatures and mixed operational characteristics.

Forecasting and planning interpretation

The regression framework was also used to project future aggregate consumption. Within this forecasting setup, the projected aggregate demand trended downward across the 2015–2025 horizon. The reported interpretation was that the household would consume approximately 8 kW less power in 2024 and about 9 kW less in 2025 relative to received power, with an average daily difference of approximately 2.5 W when comparing July 2024 and July 2025.

From an urban planning perspective, the central importance of this result is not the exact long-range value itself, but the demonstration that low-frequency NILM can support forward-looking residential energy management. When used cautiously, such predictive insights can inform user recommendations, neighbourhood demand profiling, and more efficient infrastructure planning.

DISCUSSION

The study demonstrates that a computationally modest NILM architecture can achieve meaningful appliance-level insight using only low-frequency steady-state active power. This is significant because many real-world smart meters already operate at similarly modest sampling rates. The system therefore avoids the practical barriers associated with high-frequency metering, extensive transient analysis, and large neural architectures.

The regression results show that Bagging is the most effective predictive model among the tested alternatives, offering the best balance between accuracy and robustness. The classification results likewise show that KNN is particularly well suited to the non-linear appliance separation problem posed by DRED. The optimised model's use of 21 neighbours, uniform weighting, and Euclidean distance indicates that careful tuning can improve performance without sacrificing interpretability or deployment feasibility.

At the same time, the remaining classification error reflects a core challenge in low-frequency NILM: appliances with similar steady-state active power can still be difficult to distinguish, especially when they span multiple operating categories or exhibit overlapping signatures. This limitation does not negate the utility of the method; rather, it clarifies the realistic performance ceiling of lightweight NILM under low-cost sensing conditions.

IMPLICATIONS FOR URBAN DEVELOPMENT AND SMART CITIES

The relevance of this work extends beyond household analytics. In smart-city contexts, scalable NILM can support at least four urban objectives:

1. Demand-side management: appliance-level visibility improves targeted conservation and behavioural feedback.
2. Energy affordability: residents can identify high-cost appliances, reduce unnecessary use, and manage bills more effectively.
3. Infrastructure planning: utilities gain better insight into load composition, supporting network expansion, renovation, and operational planning.
4. Sustainability and carbon reduction: improved load efficiency reduces the pressure on generation systems and contributes to lower emissions over time.

Because the method is compatible with existing low-frequency smart meters and does not require appliance-level sub-meter installation, it is particularly well aligned with resource-constrained urban deployments where cost-effectiveness and scalability are essential.

LIMITATIONS AND FUTURE RESEARCH

Several limitations define the next stage of methodological improvement.

First, the dataset is drawn from a single monitored household, which limits direct generalisation to broader housing stocks. Multi-household validation remains necessary before large-scale operational claims can be made.

Second, the DRED dataset does not include manufacturer power ratings for the appliances. This limits the ability to compare observed operating signatures against device-rated specifications, which could otherwise improve both classification and interpretability.

Third, the model is intentionally based on steady-state active power alone. While this improves deployability, it also constrains the discriminative power available for appliances with highly similar signatures or multiple operating states.

Future research should therefore incorporate appliance power ratings where available, examine the influence of external environmental conditions more systematically, and evaluate robustness across multiple datasets and household types. In addition to standard performance metrics, future model assessment should emphasise robustness, suitability for deployment, ethical considerations, and trustworthiness in real operational settings.

CONCLUSION

This study presents a practical NILM framework for low-frequency residential energy analytics based on a Bagging regressor and a modified multiclass K-Nearest Neighbour classifier. Using only steady-state active power from the DRED dataset, the framework achieves strong regression performance ($R^2 = 0.9624$) and competitive appliance classification accuracy (78.25%) while preserving computational simplicity and compatibility with conventional smart meters.

The broader significance of the work lies in its urban applicability. By enabling appliance-level insight from a single measurement point, the framework supports energy-efficient behaviour, cost control, smart-home recommendations, and more informed urban energy planning. As cities continue to pursue intelligent and sustainable infrastructure, low-cost NILM offers a credible pathway for linking household-level data to city-scale energy management outcomes.

DATA AVAILABILITY

The study is based on the Dutch Residential Energy Dataset (DRED), an open-source residential energy dataset containing aggregate and appliance-level active power measurements, together with related contextual information.

CONFLICT OF INTEREST

The authors declare no conflict of interest.

REFERENCES

[1] Hart, G. W. (1992). Nonintrusive appliance load monitoring. *Proceedings of the IEEE*, 80(12), 1870–1891.

[2] Cover, T. M., & Hart, P. E. (1967). Nearest neighbor pattern classification. *IEEE Transactions on Information Theory*, 13(1), 21–27.

[3] Breiman, L. (1996). Bagging predictors. *Machine Learning*, 24(2), 123–140.

[4] Kolter, J. Z., & Johnson, M. J. (2011). REDD: A public data set for energy disaggregation research. In *Proceedings of the SustKDD Workshop on Data Mining Applications in Sustainability*.

[5] Batra, N., Kelly, J., Parson, O., Dutta, H., Knottenbelt, W., Rogers, A., Singh, A., & Srivastava, M. B. (2014). NILMTK: An open source toolkit for non-intrusive load monitoring. In *Proceedings of the 5th ACM International Conference on Future Energy Systems* (pp. 265–276).

Adson Damasceno, Department of Technology, State University of Feira de Santana, Feira de Santana 44036-900, Brazil

Ivanovitch Silva, Applied Computing Graduation Program, State University of Feira de Santana, Feira de Santana 44036-900, Brazil

Manuscript Published; 29 August 2025.