Proceedings of the Design Society, Volume 5: ICED25
https://doi.org/10.1017/pds.2025.10163

Enabling data-driven design by deriving consumer appliance use from household energy data

Nathan Morris✉, **James Gopsill**, **Sindre Eikevåg**, **Maria Valero** and **Ben Hicks**

School of Electrical, Electronic and Mechanical Engineering, University of Bristol, UK

✉ nathan.morris.2021@bristol.ac.uk

ABSTRACT: Achieving Net Zero requires designers to have a better understanding of the product use with studies showing user behaviour, cultural norms, seasonality and product interactions concomitantly dictate energy consumption. Data on product use can support data-driven design processes that have been shown to improve the efficiency of existing products. The paper reports a method that generates data for data-driven design processes from non-intrusive load monitoring (NILM) of household energy consumption data. The method produced appliance classification accuracies of 0.9984 while reducing sample size, sampling frequency and machine learning model complexity showing potential for it to be deployed at scale across communities.

KEYWORDS: data-driven design, non-intrusive load monitoring, collaborative data, wavelets

1. Introduction

The design of energy efficient products and services is a non-trivial task. Cabeza et al. (2014) revealed products sold with improved energy efficiency resulted in more energy being consumed as users operate the products for longer periods of time. Cultural norms, such as 'getting a cup of tea' at the half-time break of a football/soccer match, and societal initiatives, such as turning on/off lights when entering rooms, drive energy use. Product interactions where the existence of another product (e.g., toasters & kettles) can also result in more energy use with users moving beyond a cup of tea to having tea and toast. This contributes to an overall increase in the usage of energy within buildings - the opposite of the desired effect of reducing energy consumption. These complexities in appliance usage necessitates a more granular (detailed) appraisal of product use and interactions between products to support data-driven design processes. It can also provide insights for other stakeholders, including electricity distribution operators.

Improved energy efficiency is critical if society is to achieve its net-zero targets. Electricity usage is expected to double by 2050 (Adam Barth and David Gonzalez, 2024). Therefore, understanding product use has never been more important to aid designers in making Net Zero design decisions. An example of different usage of appliances on a residential electricity signal is shown in Figure 1. A 'coupled' interaction can be seen, with energy use in the refrigerator, microwave and kettle increasing within the same time-frame.

Having the capability to non-intrusively observe product use would give designers the information to complete more rigorous data-driven design (Deutsch, 2015). Monitoring existing and legacy devices could provide indicative trends with regard to product usage times, duration, and energy efficiency on the micro-scale. Macro insights include product energy usage, lifetime performance and end-of-life patterns, aiding in the identification of how appliances are used over the whole scope of their life. The implementation of this technique could also identify which devices are used in tandem with each other while also identifying legacy products that might still be in use, giving further insight into usage patterns.

Non-intrusive load monitoring (NILM) is the field that considers how insights can be generated from single energy sensors, such as from a smart energy meter. It requires additional data processing steps to

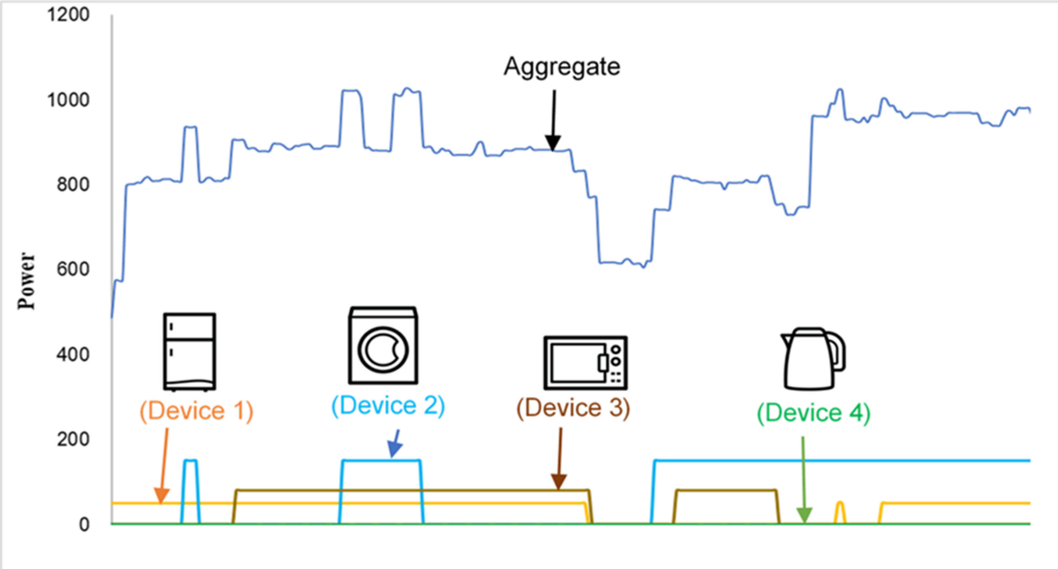

Figure 1. An example of how individual devices can be detected from an aggregated signal (Ghaffar et al., 2022)

provide contextualised breakdown of the electricity signal into its constituent components, with the main idea to produce disaggregated energy bills for consumers. NILM has seen application in residential electricity monitoring as well as other utilities such as water, gas and heat.

The main challenges of using NILM for data-driven design include:

1) Accurately identifying appliances in an ever richer transient energy signal.
2) Minimising the sampling rates required to provide accurate identification.
3) Efficient processing of data streams to attain data in a timely and sustainable manner.
4) Reporting the information back such that decisions-makers can perform decision-making.

This paper focuses on challenges 1 and 2. Current NILM has reached the limits of its accuracy, which cannot be overcome by classical techniques. The field has therefore turned to machine learning (ML) algorithms to support the classification of signals via their unique characteristics. ML algorithms can be further supported by pre-processing of the data to accentuate the features that distinguish one appliance from another. This can, in turn, reduce the energy footprint and data sampling requirements of ML workflows enabling them to be realised for real-world scenarios, such as smart energy meters. Transient signatures often have nuanced shapes, meaning high performing classification algorithms are needed alongside various signal analysis techniques. One such technique is wavelets, which can provide effective signal de-noising and characteristic enhancement, and with over 14 families the technique can be tailored for almost any application.

The contribution of this paper is the development of a wavelet pre-processing ML workflow for the disaggregation of residential electricity signals that addresses accuracy and computational effort so that the data-driven insights can be processed at scale using existing smart energy meter hardware.

The paper continues with a related work (Section 2) section that summarises previous NILM research and the application of ML algorithms and wavelets. Section 3 details the implemented workflow and its main steps, including: dataset preprocessing, wavelet transform application and ML models. Section 4 reports the results from the study, which are then discussed in Section 5. The discussion covers the feasibility of disaggregating individual appliances and the potential for data-driven design. The paper then concludes, Section 6, with the key findings and further work to implement the method.

2. Related work

This section summarises the algorithmic developments in NILM including ML applications. The section then provides a background on wavelet analysis techniques and describes its suitability as a pre-processing technique.

2.1. Non-intrusive load monitoring

NILM is the study of signal disaggregation. Example signals include electricity, gas and water readings. Signals are measured at regular intervals and the objective is to identify, categorise and characterise what devices are in use. The process is typically broken down into four steps (Hart, 1989):

1) Data collection
2) Event detection
3) Appliance classification
4) Signal disaggregation

Domestic electricity signals are typically used to test NILM methods due to availability of labelled datasets. Signals fall into two main categories: steady state and transient signatures. Steady-state signal processing was first performed by Hart (1989) using a threshold method. The method proved useful for appliances that consumed large amounts of power (150W) but was unable to capture many lower-power appliances such as Compact Fluorescent Lightbulbs. This was due to signal noise when monitored at the high frequencies required for NILM causing false positives (FP) and false negatives (FN) within the classification algorithm. This idea provides the basis for further investigation of NILM using transient signatures which are 10s of milliseconds long with more advanced algorithms, which aid to understand optimal sampling frequencies of appliance electricity usage.

2.2. Machine learning in NILM

Methods incorporating ML to help capture the subtleties in energy signal data and disaggregate lower power signals formed by a greater number of devices have been developed. Lee (2003) and Young et al. (2023) obtained F1-Scores of 0.97 and 0.989 respectively, for example.
These algorithms could streamline steps 2-4 of the NILM process, helping to generate more accurate appliance identification, that could provide greater insight to designers designing devices and/or systems of household equipment. However, the data processing, training/re-training and testing regimes of ML workflows can be highly computationally intensive. Simple models can take up to an hour to retrain and use up to 759 J of electricity (García-Martín et al., 2019). The computational intensity of existing methods also limits the potential of deploying the models on embedded devices (e.g., smart energy meters) requiring data to be sent in real-time and processed in the cloud at high data rates. Thus, the research challenge identified and explored in this paper is in optimising the ML workflow to reduce the computational burden while maintaining the accuracy required to generate insights for data-driven design processes.

2.3. Wavelets

An approach to reduce the computational burden is to accentuate the differentiating features in a data stream by pre-processing the signal. The approach examined in this paper was wavelets. Wavelets are known as “small waves” meaning they have concentrated energy in finite time (Burrus et al., 1998). They are a useful tool in signal processing due to oscillating wave-like characteristics and ability to allow simultaneous time and frequency analysis, which is especially useful for transient, time-varying phenomena. Wavelet analysis techniques use localised wavelets to give time-frequency localisation of signals that Fourier analysis cannot provide (Taspinar, 2018). They use the same underlying principles that any waveform of any shape can be deconstructed into a series of other waveforms, each with its own frequency and amplitude. In the instance of a wavelet transform, it deconstructs the original signal into a series of basic waveforms called wavelets. These transforms have successfully been applied to heart sound segmentation (HSS) methods (Milani et al., 2022) and in pattern recognition algorithms used for fault detection and diagnosis in pipes (Ruiz de la Hermosa Gonzalez-Carrato et al., 2014).
Wavelets are not simple sinusoids but complex and, at first sight, arbitrary squiggles, however they can exaggerate unique characteristics within signals. This can help the classification process used by ML models and hence identify appliances more accurately. Wavelets can be both discrete and continuous, therefore all options should be explored as part of a pre-processing step in a workflow to understand the optimal method to apply the wavelet transform for each unique application.

3. Adding wavelets to NILM ML workflows

Figure 2 details the study that was conducted to assess the benefits of applying a Wavelet pre-processing step in an NILM ML workflow. The study started by taking two existing domestic appliance signal datasets – PLAID (Medico et al., 2020) and WHITED (Kahl et al., 2016). The PLAID dataset comprises of 1876 records of individually metered appliances from 17 different appliance types comprising of 330 different makes and models. These signals were recorded in 65 different locations in Pittsburgh, Pennsylvania (US) and contain 1314 records of the combined operation of 13 of these devices too. The WHITED dataset comprises 1259 different records for 120 different appliances which can be grouped into 47 appliance types. WHITED has recorded these appliances in eight different regions across the world, meaning different electricity distribution frequencies are apparent. The datasets contain eleven common appliances with records of the electrical RMS voltage and AC current, of which the RMS current and real power were calculated.

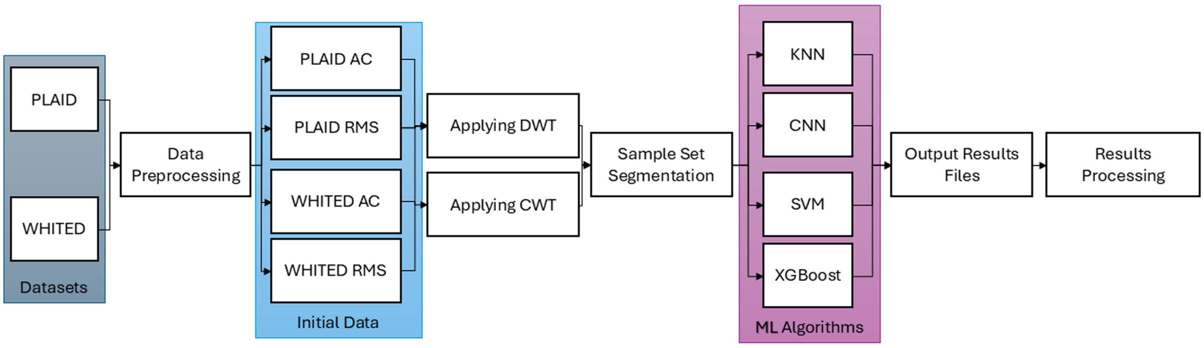

Figure 2. A diagram of the workflow used to pre-process data (light blue) from datasets (grey), apply Wavelet transforms via the CWT and DWT, execute ML algorithms (pink) and post-process results

PLAID and WHITED are both high frequency, 30 kHz and 44 kHz, respectively datasets containing eleven common appliances with records for RMS voltage and AC current, of which the RMS current and real power were calculated. The high frequency enabled us to downsample and explore a range of frequencies (10 Hz, 50 Hz, 150 Hz, 500 Hz, 1 kHz) to find a minimum frequency in order to reduce the computational overhead. These datasets were then segmented into five training/testing sets (500, 1000, 5000, 10000, 20000) to find the minimum dataset size threshold for accurate training to, again, reduce computational overhead. Event shifting and signal scaling were used to increase the number of samples of the larger segmentation sets.

The next step was to pass each of the datasets through a wavelet pre-processing step that would then result in the datasets for the ML algorithms to train and test on. Wavelet transforms were applied using the Continuous Wavelet Transform (CWT) and Discrete Wavelet Transform (DWT). The CWT was applied using the continuous Morlet wavelet and Equation (1) utilising 128 scales.

$$X_{w(a,b)} = \frac{1}{|a|^{\frac{1}{2}}} \int_{-\infty}^{\infty} x(t)\psi\left(\frac{t-b}{a}\right)\mathrm{dt} \tag{1}$$

where $x(t)$ is the original signal and $\psi(t)$ is the continuous mother wavelet which gets scaled by a factor of a and translated by a factor of b.

The DWT was applied as a filter bank and is more computationally efficient. Five discrete wavelets were chosen – Daubechies 1, Daubechies 2, Symlets 2, Symlets 4 and Biorthogonal 4.4. These were used to the 2, 7, and 12 maximum decomposition levels and selected due to their recommended usage in edge and feature detection of closely spaced features along with their denoising properties.

The output coefficients of both transform types are obtained as approximate and detailed coefficient arrays. These were then stitched together end-to-end into a one-dimensional array, downsampled back to 5 kHz and fed into the ML models for testing.

Four ML algorithms were selected to examine the benefits that the pre-processing step could bring to each one. These were the Convolutional Neural Network (CNN), K-Nearest Neighbours (KNN), Supported Vector Machine (SVM) and Extreme Gradient Boosting (XGBoost) algorithms, and have been used in the previous NILM studies.

The CNN used was configured as a standard deep CNN, which is more efficient than a dense neural network but retains complex feature information. The architecture includes an input layer, three convolution layers, a GAP layer, a dense layer and an output layer with nodes that match the number of common appliances. The KNN was selected due to its simplicity and works off the assumption that similar points can be found near each other in feature space. The SVM is similar to the KNN, but attempts to find the maximum separating hyperplane between classes. Both the KNN and SVM are adaptable and efficient, while also posses the ability to manage high-dimensional data and non-linear relationships effectively. But they both only store a training dataset rather than going through a training stage, hence they become increasingly memory intensive and inefficient as the dataset grows. The XGBoost algorithm is an implementation of the gradient boosted decision tree (GBDT) but is built in parallel rather than sequentially. This allows it to handle large datasets, be scalable and effectively handle missing values without requiring significant preprocessing.

The combination of datasets, pre-processing operations, and ML algorithms resulted in 408 scenarios to be evaluated. The datasets were split 90/10% training and testing and cross-validated using a 5-fold split resulting in 25 tests for each scenario (only 20 tests for the WHITED RMS dataset combinations). This led to a total of 9,860 runs that were performed on a single node on the University of Bristol's BlueCrystal 4 high-performance computing facility featuring 28 CPUs, two Nvidia P100 GPUs and 12GB of RAM. The results of each job were then extracted, processed and analysed.

The metrics of interest were the F1-Score (Equation (2)) against the computational overhead required to achieve the score - sample rate, dataset size and ML model. F1-score is described as a harmonic mean of the accuracy and precision of the ML models classification process. Other results of interest include data preprocessing times, model file sizes, and algorithm compute times.

$$F1 = \frac{TP}{TP + \frac{1}{2}(FP + FN)} \tag{2}$$

where TP is the number of true positives, TN is the number of true negatives, FP is the number of false Positives and FN is the number of false negatives.

An F1-Score change of 0.001 was selected as the minimum change to qualify as an accuracy improvement with a batch size of 32 and maximum number of epochs of 1500 chosen for training and testing of the CNN. The maximum number of epochs for cross-validation of the CNN was selected to be 100 due to job time constraints. It isn't necessary to specify these values for the other ML algorithms.

4. Results

The best results obtained using the basic datasets in this research were cross validated F1-Scores of 0.9954 and 0.9915 obtained with the WHITED RMS and PLAID AC combinations with the XGBoost and CNN algorithms respectively. This details the significant influence the use of WHITED RMS data has on the F1-Score obtained as an increase in F1-Score of 0.004 here in a sample size of 22000 is classifying eight appliances more accurately. This highlights the algorithm and input type combination of WHITED RMS as the best, further shown in Table 1.

Table 1 shows significantly that the WHITED RMS XGBoost combination, with a sample size of 11000, consistently produces the best F1-Scores. Sampling frequency also had an impact, with the best scores obtained from the highest frequencies of 500 Hz and 1 kHz. Wavelets also had an influence here, with Symlets 2 discrete wavelets applied using the DWT giving an F1-Score of 0.9984, the optimal combination, while the top ten cross-validated F1-Scores use the DWT. When compared with the results obtained by Young, this suggests that twenty more appliances are classified correctly when using wavelets, a significant difference when trying to accurately disaggregate electricity usage. Thus, at high fidelity levels discrete wavelets make a difference, but there was no real trend to a single wavelet or decomposition level that performs the best. It was also seen that the DWT performs better than CWT, although not by much. Figure 3 shows that the best results were obtained at the larger sample sizes and frequencies as one would expect. It also highlighted that the general trend is that the larger sample size has more impact than the higher sampling frequency. The best scores were obtained on level 2 for all frequencies and sample sizes, while level 12 is consistently better at lower frequencies than level 7 but this swaps at higher frequencies. Initial observations suggest the use of wavelets reduces the sampling frequency signals were necessary to be recorded at while still obtaining high F1-Score.

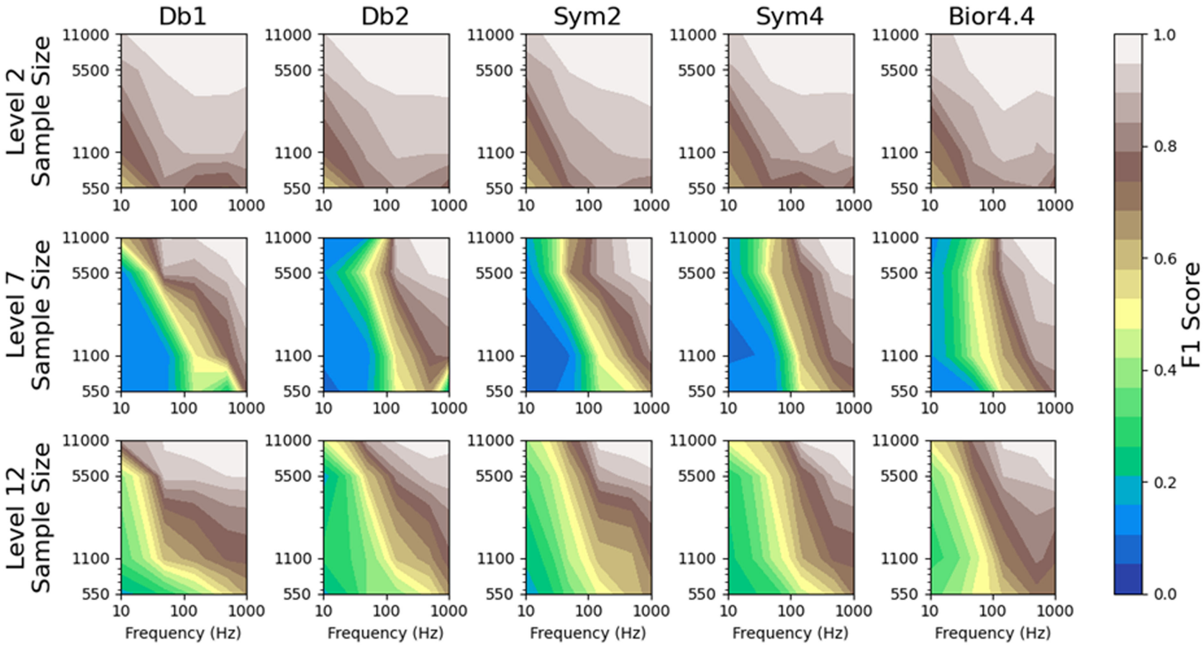

Figure 3. Heat-maps detailing the difference in F1-Scores obtained from the classification report of the CNN algorithm for each discrete wavelet combination on the WHITED RMS dataset

5. Discussion

The section offers potential reasons behind the results obtained and critically appraises the implemented workflow. It then describes the impact on key stakeholders and how collaboration is essential for the secure implementation of the presented approach and NILM technology.

5.1. Feasibility of disaggregating individual devices

ML algorithms are improving rapidly, hence, when compared to the results obtained by Young et al. (2023), the ML models perform better even with only the basic data (three package releases of TensorFlow). This also has a knock-on effect with the comparison of the updated workflow with Young's F1-Scores. The application of discrete wavelets using the DWT has a distinct impact on the

Table 1. The top five cross-validated F1-Scores obtained from the test cases with associated combination metadata

Wavelet	Level	Dataset	Input Type	Algorithm	Sample Size	Frequency	F1-Score Classification	Cross-Validated	Compute Time [s]	File Size [MB]
Sym2	2	whited	RMS	XGBoost	11000	1000	0.9961	0.9984	3460.78	1.1783
Db1	7	whited	RMS	XGBoost	11000	500	0.9935	0.9980	1735.32	1.2177
Sym4	2	whited	RMS	XGBoost	11000	1000	0.9967	0.9979	3451.98	1.1790
Db1	7	whited	RMS	XGBoost	11000	1000	0.9951	0.9976	3211.28	1.1359
Bior4.4	12	whited	RMS	XGBoost	11000	1000	0.9972	0.9976	3173.88	1.2312

F1-Scores compared to the CWT application, which was most likely caused by arrays that are too disjunct. The CWT outputs large amounts of coefficients compared to the DWT. These were then downsampled to 5 kHz to facilitate the saving of data, which was likely to lose large amounts of the characteristic feature data that the CWT amplifies. A similar pattern was most likely seen with the DWT at higher decomposition levels. In effect, too much data that isn't continuous can cause issues when inputted into ML models, resulting in the lower levels of decomposition being seen as optimal, while there is also no obvious discrete wavelet that performs exceptionally well.

The result of a cross-validated F1-Score of 0.9984 suggests that a significant part of solving challenge 1, accurately identifying appliances, has been met and is critically consistent. The following result of a 0.9980 F1-Score on a 500 Hz frequency combination and Figure 3 also demonstrated minimal sampling rates still provide consistent accurate appliance classification, laying the foundations for challenge 2 to be further explored.

A key point of this research was to examine the data generation and ML model compute times, as well as the dataset and model file sizes. Generally, dataset generation times were not excessive, with the longest of the combinations being 88 s for 80,000 labelled samples on the PLAID RMS side. RMS data generation took longer than AC for both the PLAID and WHITED, which was likely due to the RMS calculation that must take place from the initial data. The DWT application time is the shortest at approximately 0.12 ms per sample, while the CWT requires 98.1 ms due to the large number of coefficients produced. The DWT application has negligible effect on the dataset file sizes, which reached a maximum of 3.2GB for the PLAID AC set, while the CWT required downsampling due to the dataset being too large to be saved by NumPy. This distinguished the DWT as the optimal wavelet transform, but also highlighted the importance of keeping testing and training set sizes to a minimum. Each training sample of a frequency of 5 kHz has a file size of approximately, 20 kB and therefore transmission of at least 10,000 samples would require 12Wh of electricity. This would cost approximately £0.003 and produce 1.7 g of CO_2 per transmission, which for every household in a country could amount to large energy requirements e.g. 336 MWh, assuming 28 Million UK households.

The model compute times on the other hand display an interesting result of how much increasing the frequency effects the compute times of models. This doubles when the frequency doubles, as would be expected due to twice the amount of data passing through the models, with a similar effect seen when increasing the sample sizes. In general, the KNN runs the fastest followed by the XGBoost and CNN algorithms while the SVM runs the slowest. ML model file sizes are very dependent on the algorithm with the CNN unchanging at 3.114MB for every combination while the XGBoost reaches a maximum of 1.23MB. XGBoost file sizes are heavily dependent on the combination present but are generally 40% the size of those generated by the CNN. The file sizes of the KNN and SVM are generally much larger and smaller respectively than the CNN and XGBoost algorithms. These factors are a critical part of integration into embedded systems and is the primary reason a trade-off study to understand the optimal combination of the workflow is necessary.

One of the key assumptions used in NILM research is identified by Hart (1992), “in a small time interval it is expected a small number of appliances change state in a typical load” and that “operating power is never negative”. This doesn’t have any impact in this isolated research but would have significant impact during integration into real-world systems. Similarly, the use of one-second windowing is likely to cause issues as some appliances’ distinct transient signals only show after a second after startup. The concatenation of the output coefficients of the DWT and CWT end-to-end is also a criticism of the approach, as this can feed discontinuities and therefore noise into the signals. When the DWT is applied at higher levels, there is also the possibility that noise is fed into the system, as too many decomposition levels can be applied to a signal. However, it is expected that the ML models are sophisticated enough to deal with these imperfections in the input data.

5.2. Potential usage information

NILM is predicted to provide electricity savings of up to 12% to the average consumer if implemented correctly (Gopinath et al., 2020). Accurate insights into electricity consumption in a domestic setting and the context around product use is of potential benefit to three main stakeholders: the consumer, the product designer and energy distribution operators. NILM can provide valuable insights into consumer behaviour and preferences by uncovering patterns of appliance usage not available by other means. This includes duration and timing of use, and usage correlation to other appliances. It could also provide information about a product’s life cycles, user habits and uncover demographic patterns – regions and socio-economic groups. It might for example, reveal how early adopters who are open to trying new items show different behaviours to laggards that may require additional incentives or education to accept new products and use them. Given this, it could help track which of these incentives are most beneficial, while also tracking life duration of the product before it is replaced by different demographics. Further, it could provide information on the macro context of energy usage while also providing micro usage detail, with the idea of disaggregated household electricity bills that would provide a detailed breakdown for the consumer that could influence their usage patterns.

5.3. Informing data-driven design processes

Many design processes are implementing data-driven methods and techniques to support future design iterations. Information on device use can aid designers in their decision-making by enabling them to evidence the need for new features that will support their customer use cases. These insights enable designers to create user-centric product designs that align closer with evolving user needs, environmental goals and actual usage. It can inform decisions about how to implement energy-saving features, device usability and product durability, while helping to identify inefficiencies and opportunities for improvement. Examples of how this data can be used to inform product feature development include the timing of energy-intensive cycles, predictive maintenance or optimisation for specific usage scenarios. Further, consumer usage information can help design products that proactively and reactively change based on user behaviours, including automatic adjustments or alerts to encourage energy efficient usage.

5.4. Enabling collaboration

To leverage NILM models, energy providers, designers, manufacturers and consumers need to work together to provide a streamlined, secure and data-sensitive approach. Manufacturers would need to provide signatures of current and legacy products, while energy providers and consumers would be essential in capturing real time data. This means providers need to report or provide access to usage information securely, or via a third party (intermediary) that could report and control access to the data for each stakeholder/group. Each stakeholder has a different interest in this technology with energy providers looking for improvement of grid management, targeted energy-saving advice and optimised pricing models. In contrast, manufacturers and designs are looking for data to drive design and refinement of products. These stakeholders could potentially use this information to offer tailored product features or services to different user segments. Users on the other hand are interested in understanding their disaggregated electricity usage and opportunities to save on utility bills wherever possible. Lastly, with any collaborative endeavour, issues of data access, trust and privacy will need to be addressed. There would also be legal and organisation hurdles to ensure smooth flow of information between industries if this technology is implemented. But it can also be used to support new legislation, user education and design interventions, aiding the drive to more efficient energy usage and helping to hold manufacturers and designers accountable to aspire for improved product electricity usage.

6. Conclusions

This paper presented the need for detailed user consumption information of appliances for future datadriven design process and that NILM is a method of generating such data. It provided a background into the previous work that has been completed, the technological barriers that presently limit the accuracy and practicality of these models and how ML models and wavelet transforms can be applied to maintain detection accuracy at low sampling rates, meaning that the approach can be implemented into existing domestic household supply. Various datasets and preprocessing steps were executed before wavelet transforms were applied and the ML models run with 9,860 unique testing scenarios. The main trends identified from the testing scenarios were:

- The WHITED RMS 1kHz sampling frequency and 11,000 sample size combination was the most accurate.
- DWT performed better than the CWT.
- The top cross-validated F1-Score obtained was 0.9984 by the Symlets 2 discrete wavelet applied to two decomposition levels.
- The optimal combination for accurate classification while considering the computational overhead was Daubechies 1 WHITED RMS XGBoost combination with a sample size of 11,000 and sampling frequency of 500Hz, which achieved a score of 0.9980.

The combination of wavelets and ML enables a reduction in the sampling frequency required making the workflow more viable than without. This is essential for collecting information to support data-driven design, as the lower sampling frequency required makes for easier integration onto embedded hardware systems and current domestic supply. Given this possibility the potential applications of the data and insights were discussed with particular attention on design and designers. This included enabling designers to create user-centric product designs that align closer with evolving user needs, environmental goals and actual usage (data that is not currently possible to acquire on mass). It could also inform decisions about how to implement energy-saving features, device usability and product durability, while helping to identify inefficiencies and opportunities for improvement, including for example, reduced efficiency through-life. It is because of these potential benefits that NILM is cited as offering the potential to reduce consumer electricity usage by 12% (Gopinath et al., 2020) and a key enabler for achieving Net Zero.

Acknowledgements

Support was provided by the Engineering Systems, Design and Innovation group and the BlueCrystal 4 HPC facilities team at the University of Bristol.

References

Adam Barth and David Gonzalez (2024). *Integrating renewable energy sources into grids | McKinsey*.

Burrus, C. S., Gopinath, R., and Guo, H. (1998). *Introduction to wavelets and wavelet transforms : a primer*. Prentice Hall, Upper Saddle River, N.J. Type: Book.

Cabeza, L. F., Urge-Vorsatz, D., McNeil, M. A., Barreneche, C., and Serrano, S. (2014). Investigating greenhouse challenge from growing trends of electricity consumption through home appliances in buildings. *Renewable and Sustainable Energy Reviews*, 36:188–193.

Deutsch, R. (2015). *Data-Driven Design and Construction: 25 Strategies for Capturing, Analyzing and Applying Building Data*. John Wiley & Sons. Google-Books-ID: MPtcCgAAQBAJ.

García-Martín, E., Rodrigues, C. F., Riley, G., and Grahn, H. (2019). Estimation of energy consumption in machine learning. *Journal of Parallel and Distributed Computing*, 134:75–88.

Ghaffar, M., Sheikh, S., Naseer, N., Mohy-Ud-Din, Z., Rehman, H., and Naved, M. (2022). Non-Intrusive Load Monitoring of Buildings Using Spectral Clustering. *Sensors*, 22:4036.

Gopinath, R., Kumar, M., Prakash Chandra Joshua, C., and Srinivas, K. (2020). Energy management using nonintrusive load monitoring techniques – State-of-the-art and future research directions. *Sustainable Cities and Society*, 62:102411.

Hart, G. (1989). Residential energy monitoring and computerized surveillance via utility power flows. *IEEE Technology and Society Magazine*, 8(2):12–16. Type: Journal Article.

Hart, G. (1992). Nonintrusive appliance load monitoring. *Proceedings of the IEEE*, 80(12):1870–1891. Type: Journal Article.

Kahl, M., Haq, A., Kriechbaumer, T., and Jacobsen, H.-A. (2016). WHITED - AWorldwide Household and Industry Transient Energy Data Set. Type: Dataset.

Lee, K. D. (2003). *Electric load information system based on non-intrusive power monitoring*. PhD Thesis, Massachusetts Institute of Technology.

Medico, R., De Baets, L., Gao, J., Giri, S., Kara, E., Dhaene, T., Develder, C., Bergés, M., and Deschrijver, D. (2020). *A voltage and current measurement dataset for plug load appliance identification in households*. ISBN: 2052-4463 Issue: 1 Pages: 49 Type: Dataset Volume: 7.

Milani, M. G. M., Abas, P. E., and De Silva, L. C. (2022). A critical review of heart sound signal segmentation algorithms. *Smart Health*, 24:100283.

Ruiz de la Hermosa Gonzalez-Carrato, R., García Márquez, F. P., Dimlaye, V., and Ruiz-Hernández, D. (2014). Pattern recognition by wavelet transforms using macro fibre composites transducers. *Mechanical Systems and Signal Processing*, 48(1):339–350.

Taspinar, A. (2018). *A guide for using the Wavelet Transform in Machine Learning*.

Young, T. L., Hicks, B., Gopsill, J., and Valero, M. (2023). A Trade-Off Study in Applying Machine Learning Algorithms to Non-Intrusive Load Monitoring. *Energy and AI*. Type: Journal Article.