

An End-to-End Machine Learning based Unified Architecture for Non-Intrusive Load Monitoring

Syed Wali

Faculty of Electrical and Computer Engineering
NED University of Engineering and Technology
Karachi, Pakistan
Syedwali110@hotmail.com

Muhammad Hassan ul Haq

Faculty of Electrical and Computer Engineering
NED University of Engineering and Technology
Karachi, Pakistan
hassanulhaq@neduet.edu.pk

Majida Kazmi

Faculty of Electrical and Computer Engineering
NED University of Engineering & Technology
Karachi, Pakistan
majidakazmi@neduet.edu.pk

Saad Ahmed Qazi

Neurocomputation Lab, National Centre of Artificial
Intelligence, NED University of Engineering and
Technology, Karachi, Pakistan
saadqazi@neduet.edu.pk

Abstract-Non-Intrusive Load Monitoring (NILM) or load disaggregation aims to analyze power consumption by decomposing the energy measured at the aggregate level into constituent appliances level. The conventional load disaggregation framework consists of signal processing and machine learning-based pipelined architectures, respectively for explicit feature extraction and decision making. Manual feature selection in such load disaggregation frameworks leads to biased decisions that eventually reduce system performance. This paper presents an efficient End-to-End (E2E) approach-based unified architecture using Gated Recurrent Units (GRU) for NILM. The proposed approach eliminates explicit feature engineering and has a unified classification and prediction model for appliance power. This eventually reduces the computational cost and enhances response time. The performance of the proposed system is compared with conventional algorithms' with the use of recall, precision, accuracy, F1 score, the relative error in total energy and Mean Absolute Error (MAE). These evaluation metrics are calculated on the power consumption of top priority appliances of Reference Energy Disaggregation Dataset (REDD). The proposed architecture with an overall accuracy of 91.2 and MAE of 25.23 outperforms conventional methods for all electrical appliances. It has been showcased through a series of experiments that feature extraction and event-based approaches for NILM can readily be replaced with E2E deep learning techniques allowing simpler and cost-efficient implementation pathways.

Keywords-non-intrusive load monitoring; gated recurrent units; end-to-end machine learning; reference energy disaggregation dataset

I. INTRODUCTION

Energy demand is increasing drastically with the increase in industrial development. This raises the need of managing energy usage effectively at consumer end. Efficient demand management is possible by analyzing the appliance level power consumption in buildings [1]. Today, this financially feasible solution, provided in 1992 [2], is known as Non-Intrusive Load

Monitoring (NILM). The basic idea of NILM revolves around the decomposition of total demand into appliance level power consumption. Since this load disaggregation approach does not depend on several data recorder sensors, it is therefore a cost-effective solution adopted for demand reduction and load forecasting. The interest of researchers in this domain is significantly increasing with the development of smart meters, capable of delivering aggregate power information to the customer. With the advancements in Machine Learning (ML) domain, it is expected that NILM with high accurate power consumption analysis capabilities will serve as the backbone in the development of innovative smart grid services [3].

Generally, NILM can be categorized into two types: event-based approaches and event-less or state dependent approaches. In the event-based approaches, any significant change in the signal that is considered during load disaggregation is regarded as an event. All event-based approaches depend on previous training, thus supervised ML approaches are mostly adopted in this category [4]. The second category, i.e. the event-less approach, does not rely on event detection. It primarily uses statistical and probability-models to match consumption signal of single or group appliances to the aggregate power signal [3]. Thus, label transitions are not required in this category. Event-based NILM methods primarily depend on finding edges in order to observe change in power demand. Initially, Hart formulated different clusters based on similar power change and characteristics of appliances using Combinatorial Optimization (CO) for disaggregating power demand [2]. Conventional edge detection approaches were then replaced by probability-based methods. These methods were comparatively less complex than the conventional methods [5, 6]. Standard deviation was calculated in [5] instead of using a single fixed parameter. Also, authors in [6] utilized a statistical approach for detecting the edges and power change signature of different appliances. Later, classifiers such as Support Vector Machines (SVMs) [7], Decision Tree [8], and other hybrid approaches

Corresponding author: Majida Kazmi

were investigated to serve the purpose. Hidden Markov Models (HMMs) and their several variants were also utilized to model multi-state appliances and different possibilities of their combination [9-11], since the complexity of modeling multistate appliances increases with increasing number of appliances installed in customer premises [12]. Therefore, inherent complexity of this methodology was reduced via the introduction of Viterbi algorithm in [13].

The latest development in ML shifts the paradigm to deep learning based NILM. Multilayer Perceptron (MLP), Convolution Neural Networks (CNNs), Deep Neural Networks (DNNs), Recurrent Neural Networks (RNN) K-Nearest Neighbor (k-NN) and Naïve Bayes classifiers are a few of the most widely used supervised ML techniques for load disaggregation [3]. CNNs and RNN-based Long Short-Term Memory (LSTM) have been explored for NILM in [12]. Authors in [14] implemented 3 different DNN architectures for short term load forecasting. 1D CNN was implemented in [15] for examining the effect of variables dependent on power demand. Hybrid CNN was also proposed in [16] where impact of reactive power, current, and apparent power were assessed on the performance of NILM disaggregation. The window of aggregate power signal was utilized in [15] to predict the power of the targeted appliance. The input sequence utilized for generating the sequence of output power is termed as sequence-to-sequence approach. If the same sequence predicts power at specific time instant only then it is termed as sequence-to-point. Both sequence-to-point and sequence-to-sequence NILM approaches [15] were based on CNN architecture. The load disaggregation problem was treated as noise reduction in [17] using denoising autoencoders. This approach showed improved performance under different types of loads.

One of the major drawbacks in the above stated deep learning based NILM architectures is their dependency on explicit features extraction from signal. This manual feature(s) extraction leads to biased decisions that eventually reduce the overall performance of NILM. In order to improve performance, these methods deploy extremely dense neural architecture with large number of layers. These dense architectures are time consuming and computationally expensive. Moreover, dependency on separate classification and regression networks [15, 18] further increases computational power, thus making these solutions extremely expensive. The manhours required in feature engineering and in the collection of contextual information for deep neural architectures comprise a time taking procedure and there lies a strong probability of losing important load signatures in manual feature extraction. In order to address all of the above limitations in previously proposed NILM approaches, this paper presents an efficient end-to-end (E2E) ML based unified architecture using Gated Recurrent Units (GRU).

The main characteristics of the proposed architecture are:

- The E2E ML approach is adopted which does not depend on explicit feature extraction. The feed input of this E2E architecture is complete aggregate power signal, ensuring reliable and better prediction even under different load categories.

- A unified module is an inherent characteristic of the proposed E2E architecture. Since the proposed architecture considers both classification of appliance and prediction of consumption as a single problem, there is no need to use separate modules.
- Low computational cost due to the unified architecture as compared to the conventional pipelined architecture. Real time load disaggregation is also possible due to the fast response time of the proposed architecture. It allows easy integration in modern smart metering devices.
- Improved performance of the proposed E2E architecture as compared to previously proposed DNN architectures despite of using comparatively lesser number of layers and neurons. The performance edge of the proposed approach is showcased on REDD, which is a renowned load disaggregation dataset.

II. THE PROPOSED E2E MACHINE LEARNING BASED UNIFIED ARCHITECTURE

An E2E ML based unified architecture for NILM is proposed in this work to completely eliminate the reliance on feature extraction. The proposed framework is presented in Figure 1. It consists of dataset and preprocessing, E2E ML model, and evaluation metrics.



Fig. 1. The proposed E2E ML-based unified architecture for NILM

A. Load Disaggregation Dataset and Preprocessing

Training and development of E2E ML Model depends on datasets prepared for load disaggregation. The Residential Energy Disaggregation Dataset (REDD) [4] is utilized in this research work. REDD was made publicly available in 2011 [4] with the aim of fast paced research and development in load disaggregation domain. It was developed with two major objectives. Firstly, it helps the researchers to apply algorithms and techniques directly on the available data instead of investing extensive efforts on the data acquisition stage of the NILM. Secondly, this dataset provides globally accepted reference data for comparing different algorithms and techniques implemented by different researchers.

The REDD dataset contains information of the aggregate power signals and the power of each individual appliance installed in 6 different homes in Massachusetts, USA. These 6 different homes cover almost all types of appliances used in consumer premises like washing machine, microwave, fridge, lights, air conditioning, electric stove, smoke detectors, etc. This dataset includes two-state, finite state, and continuously varying type of electrical loads. Low frequency power data from the first building of REDD dataset are selected for evaluation of the proposed algorithm. Six top priority appliances with respect to power consumption from House1 of REDD were considered. The selected dataset is preprocessed for removing erroneous readings, detecting gaps and downtime

using NILM tool kit (NILMTK). The data are separated as training and testing subsets with a ratio of 50:50

B. The E2E ML Model

The Gated Recurrent Unit (GRU), a variant of RNN, is selected as the basic ML model for the proposed E2E architecture. RNN is selected due to its efficient information handling with smaller context [19]. GRU is computationally simpler as compared to other RNN variants. It controls the flow of contextual information using just two gates as illustrated in Figure 2.

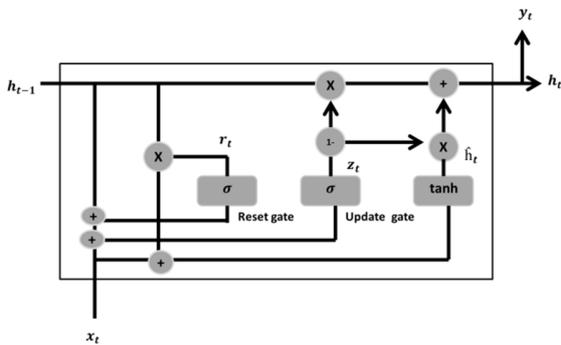


Fig. 2. The GRU architecture.

The first gate of GRU is the update gate. This gate is responsible to decide the extent to which information content from previous history should be passed for the determination of the future state. The output vector of the update gate (z_t) depends on previous cell output (h_{t-1}), present input (x_t), calculated weights (W_z, U_z), and biased vector (b_z). Mathematically, the output of the update gate depends on the sigmoid function and can be represented as:

$$z_t = \sigma [W_z * X_t + U_z * h_{t-1} + b_z] \quad (1)$$

The second GRU gate is the forget gate. This gate is responsible to filter and remove the flow of information from cells. It depends on current input (x_t), previous output (h_{t-1}), corresponding weights (W_r, U_r) and biased vector (b_r):

$$r_t = \sigma [W_r * X_t + U_r * h_{t-1} + b_r] \quad (2)$$

The final output produced by the GRU depends on intermediate memory state (\hat{h}_t). This intermediate memory depends on weights (W_h, U_h), current input (x_t), previous output (h_{t-1}), and biased vector (b_h). The mathematical model of this hidden memory state is shown in (3). It depends on tanh which is used as the activation function.

$$\hat{h}_t = \tanh [W_h * X_t + r_t * U_h * h_{t-1} + b_h] \quad (3)$$

The GRU output depends on the hidden memory state and update gate as shown in (4):

$$h_t = (1 - z_t) * h_{t-1} + z_t * \hat{h}_t \quad (4)$$

The proposed deep neural architecture consists GRU hidden layers, whereas the convolution layer and the dense layer with linear activation function are the input and output layers. The number of layers and neurons for GRU is selected for optimal performance using constructive approach in multiple passes. It

starts with a small or undersized network having a smaller number of neurons and layers, and gradually its size increases until it achieves optimal performance. During the first pass, the number of neurons is steadily increased from 64 to 2100 with a step size of 100. Better performance in terms of accuracy and MAE is observed in the range of 500 to 700 neurons as shown in Figure 3. In the second pass, neurons are gradually increased from 500 to 700 with a reduced step size of 10. It was found that the load disaggregation model suffered from overfitting when the number of neurons increased above 650. The best point in this region in terms of reduced MAE and improved accuracy was 630 neurons, so it was selected as the optimum number of neurons.

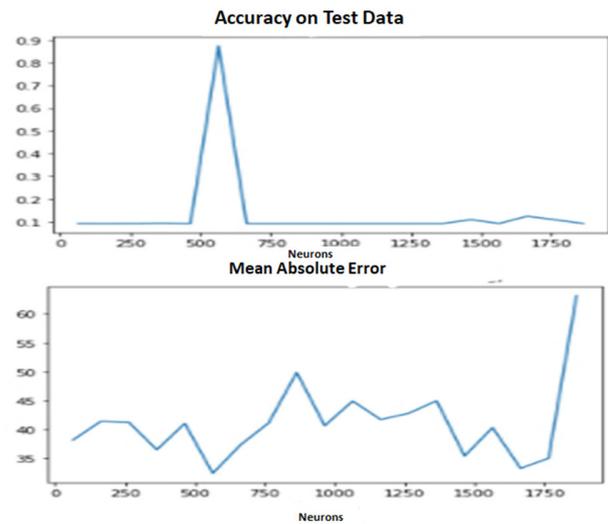


Fig. 3. GRU accuracy and MAE for different numbers of neurons.

Single hidden layered architecture is insufficient as a network with a smaller number of layers or neurons often fails to extract details from the training data [20]. Thus, 6 different architectures were tested comprising of 2, 4, 6, 8, 10, and 12 hidden GRU layers. Increasing architecture after 4 layers leads to overfitting and loss of generalization. Table I tabulates the accuracy on training data against the number of GRU layers. Optimal performance was achieved with the model with 4 GRU hidden layers of 1, 630, 1, and 1 neurons respectively.

TABLE I. ACCURACY ON TRAINING DATA AGAINST THE NUMBER OF GRU LAYERS

No. of layers	Accuracy						
	Overall test data	Fridge	Microwave	Washer dryer	Dish washer	Light	Socket
2	0.86	0.82	0.92	0.88	0.83	0.72	1
4	0.95	0.97	0.95	0.98	0.89	0.91	1
6	0.87	0.92	0.91	0.91	0.78	0.67	1
8	0.85	0.88	0.91	0.81	0.79	0.73	1
10	0.82	0.79	0.88	0.77	0.75	0.73	1
12	0.80	0.77	0.86	0.75	0.71	0.72	1

The learned model was then applied to the testing data. Table II tabulates the accuracy on the test data. The training and testing phases of the proposed E2E ML model are elaborated using the pseudocode shown in Figure 4.

TABLE II. ACCURACY OF THE LEARNED MODEL ON TEST DATA

Appliance	Accuracy
Fridge	0.88
Microwave	0.93
Washer dryer	0.98
Dish washer	0.81
Light	0.87
Socket	1
Overall	0.912

Algorithm: Pseudocode for training and testing of proposed

```

Import libraries
Import REDD dataset
Select House 1
Train Data ← Set train window in REDD
Test Data ← Set test window in REDD
Loads=[top appliances in House 1]
Algorithms = Proposed E2E Model
Initialize Arrays & Dataframe for storing results
For A in Algorithms:
    Initialize Disaggregator( ) w.r.t A
    For L in Loads:
        Train Disaggregator(Train data, L, sample period)
    end
    For L in Loads:
        R=Test Disaggregator( Test data, sample period )
        Results.append (R)
        Calculate True Positive, False Positive, True Negative
        Calculate Recall, Precision, Accuracy, F1 Score, MAE
        Store Accuracy metrics in Array
    end
    Pandas Dataframe ← Appliance Results
end
Display Results of Proposed E2E Model stored in Dataframe
    
```

Fig. 4. Training and testing pseudocode of the proposed E2E ML model.

C. Evaluation Metrics

The prediction of each load consumption at a certain time instant is not purely a regression problem. The load disaggregator first identifies the presence of certain appliances

in the aggregate power signal and then predicts their individual power consumption. Thus, performance evaluation should not be based on regression indices only, classification accuracy must also be evaluated. Recall, precision, accuracy, and F1 score are used to account classification performance, whereas, Mean Absolute Error (MAE) and relative error (RE) in total energy are used for the assessment of values predicted by the load disaggregator:

$$\text{Recall} = \frac{\text{True Positive}}{\text{False Negative} + \text{True Positive}} \quad (5)$$

$$\text{Precision} = \frac{\text{True Positive}}{\text{False Positive} + \text{True Positive}} \quad (6)$$

$$\text{F1 score} = 2 * \frac{\text{Precision} * \text{recall}}{\text{Precision} + \text{recall}} \quad (7)$$

$$\text{Accuracy} = \frac{\text{True Positive} + \text{True Negative}}{\text{Ground Positive} + \text{Ground Negative}} \quad (8)$$

$$\text{RE} = \frac{|\text{Total predicted energy} - \text{Total actual energy}|}{\max(\text{Predicted Energy}, \text{Actual Energy})} \quad (9)$$

$$\text{MAE} = \frac{\sum_{t=1}^{T} |\text{estimated power at } t - \text{actual power at } t|}{\text{Time}(T)} \quad (10)$$

III. RESULTS AND COMPARISON

The proposed E2E ML model was evaluated and compared with conventional load disaggregation algorithms. NILMTK is equipped with benchmark algorithms. Two of the most widely used benchmark algorithms, i.e. CO and FHMM, are selected for evaluation on the REDD dataset, whereas one state of the art algorithm based on RNN architecture proposed in [21] is reproduced using the details provided by the authors. For fair comparison, all baseline algorithms were trained and tested with similar train tests split on 6 top priority appliances of REDD, i.e. fridge, microwave, dishwasher, washer dryer, sockets, and lights. These top 6 appliances, with respect to power consumption, were used to evaluate the performance of the conventional and the proposed algorithms as shown in Figure 5.

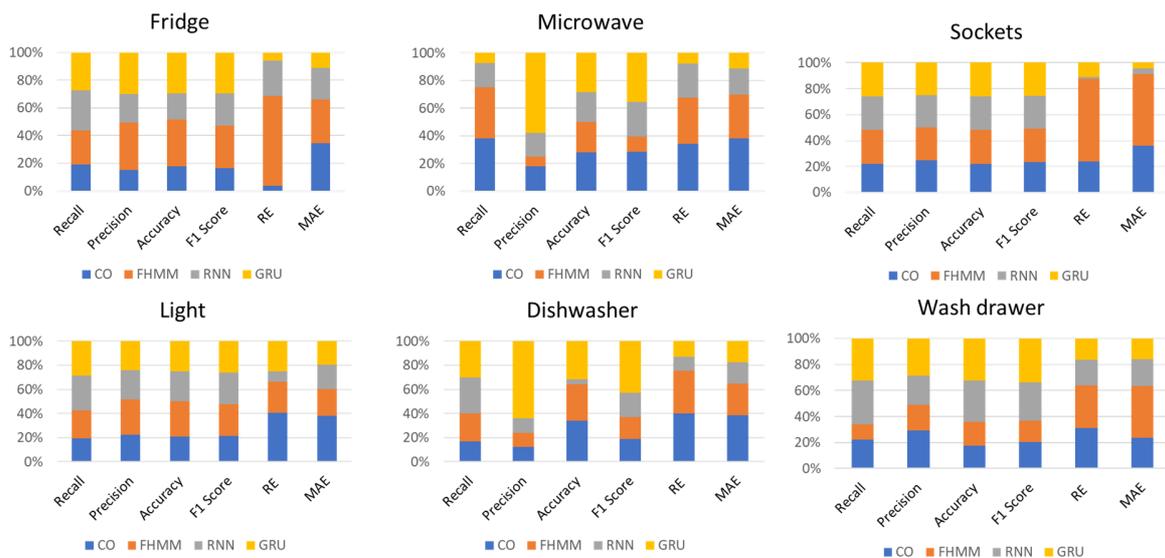


Fig. 5. Performance Evaluation of the NILM algorithms.

The proposed GRU shows a consistent performance in all the appliances, in terms of MAE. The overall MAE of the proposed model is 25.23, whereas, the MAE of CO, FHMM, and RNN is 61.88, 56.14, and 33.92 respectively. Thus, the proposed approach shows reduced MAE without being dependent on feature engineering approach. Similarly, the overall accuracy of CO, FHMM, RNN, and the proposed model is calculated as 66, 74, 64, and 91 respectively. FHMM shows better accuracy in fridge and light only, whereas, CO is better in the case of the dishwasher. This clearly indicates a better and consistent performance of the proposed model in terms of accuracy.

Beside baseline algorithms, the proposed model is also compared with few of the latest DNN approaches Convolution Sequence to point (Seq2point) [15], Convolution sequence to Sequence (Seq2Seq) [15], and Denoising Autoencoder (DAE) [17]. These DNN approaches are compared in terms of MAE and computational cost for two major appliances (fridge and microwave) of REDD's House1. The number of layers depicts the computational complexity and time required for power disaggregation. Table III indicates that the proposed E2E architecture shows better performance in comparison to the modern approaches in [15, 17], whereas this unified and less dense architecture is also computationally efficient and fast to be used in real time applications.

TABLE III. COMPARISON WITH LATEST DNN APPROACHES

Ref.	Algorithm	Layers	Comp. cost	Fridge MAE	Microwave MAE
[15]	Seq2Point	8	Medium	28.10	28.19
[15]	Seq2Seq	8	Medium	30.63	33.27
[17]	DAE	10	Highest	29.11	23.26
Proposed	E2E	6	lowest	22.12	23.19

IV. CONCLUSION

NILM is dominated by pipeline approaches with explicit feature extraction for load identification. The developments made in the domain of ML can eliminate the dependency on inherently low performing and computationally complex architectures. The proposed E2E ML based unified architecture using GRUs was evaluated on top 6 appliances of REDD. Comparative analysis of the proposed and previously reported algorithms showed that the proposed architecture has the ability to replace the feature-based load disaggregation approach. Since the proposed architecture performs better in all cases irrespective of the appliance type, it depicts that it can improve performance of the load disaggregator and due to its less computational demand and fast response time it can also be integrated in modern smart metering solutions.

REFERENCES

[1] M. Sun, F. M. Nakoty, Q. Liu, X. Liu, Y. Yang, and T. Shen, "Non-Intrusive Load Monitoring System Framework and Load Disaggregation Algorithms: A Survey," in *2019 International Conference on Advanced Mechatronic Systems (ICAMechS)*, Kusatsu, Japan, Aug. 2019, pp. 284–288, <https://doi.org/10.1109/ICAMechS.2019.8861646>.

[2] G. W. Hart, "Nonintrusive appliance load monitoring," *Proceedings of the IEEE*, vol. 80, no. 12, pp. 1870–1891, Dec. 1992, <https://doi.org/10.1109/5.192069>.

[3] L. Pereira and N. Nunes, "Performance evaluation in non-intrusive load monitoring: Datasets, metrics, and tools—A review," *WIREs Data Mining and Knowledge Discovery*, vol. 8, no. 6, 2018, Art. no. e1265, <https://doi.org/10.1002/widm.1265>.

[4] J. Kolter and M. Johnson, "REDD: A Public Data Set for Energy Disaggregation Research," *Artificial Intelligence*, vol. 25, pp. 59–62, Jan. 2011.

[5] M. E. Berges, E. Goldman, H. S. Matthews, and L. Soibelman, "Enhancing Electricity Audits in Residential Buildings with Nonintrusive Load Monitoring," *Journal of Industrial Ecology*, vol. 14, no. 5, pp. 844–858, 2010, <https://doi.org/10.1111/j.1530-9290.2010.00280.x>.

[6] D. Luo, L. K. Norwood, S. R. Shaw, and S. Leeb, "Monitoring HVAC Equipment Electrical Loads from a Centralized Location: Methods and Field Test Results," *ASHRAE Transactions* 2002, vol. 108, no. 1, pp. 841–857, 2002.

[7] H. Altrabalsi, V. Stankovic, J. Liao, L. Stankovic, and Department of Electronic and Electrical Engineering, University of Strathclyde, Glasgow, UK, "Low-complexity energy disaggregation using appliance load modelling," *AIMS Energy*, vol. 4, no. 1, pp. 1–21, 2016, <https://doi.org/10.3934/energy.2016.1.1>.

[8] J. Liao, G. Elafoudi, L. Stankovic, and V. Stankovic, "Non-intrusive appliance load monitoring using low-resolution smart meter data," in *2014 IEEE International Conference on Smart Grid Communications (SmartGridComm)*, Venice, Italy, Nov. 2014, pp. 535–540, <https://doi.org/10.1109/SmartGridComm.2014.7007702>.

[9] W. Kong, Z. Y. Dong, D. J. Hill, J. Ma, J. H. Zhao, and F. J. Luo, "A Hierarchical Hidden Markov Model Framework for Home Appliance Modeling," *IEEE Transactions on Smart Grid*, vol. 9, no. 4, pp. 3079–3090, Jul. 2018, <https://doi.org/10.1109/TSG.2016.2626389>.

[10] D. Bajović, K. He, L. Stanković, D. Vukobratović, and V. Stanković, "Optimal Detection and Error Exponents for Hidden Semi-Markov Models," *IEEE Journal of Selected Topics in Signal Processing*, vol. 12, no. 5, pp. 1077–1092, Oct. 2018, <https://doi.org/10.1109/JSTSP.2018.2851506>.

[11] R. Bonfigli, E. Principi, M. Fagiani, M. Severini, S. Squartini, and F. Piazza, "Non-intrusive load monitoring by using active and reactive power in additive Factorial Hidden Markov Models," *Applied Energy*, vol. 208, pp. 1590–1607, Dec. 2017, <https://doi.org/10.1016/j.apenergy.2017.08.203>.

[12] L. Mauch and B. Yang, "A new approach for supervised power disaggregation by using a deep recurrent LSTM network," in *2015 IEEE Global Conference on Signal and Information Processing (GlobalSIP)*, Orlando, FL, USA, Dec. 2015, pp. 63–67, <https://doi.org/10.1109/GlobalSIP.2015.7418157>.

[13] S. Makonin, F. Popowich, I. V. Bajić, B. Gill, and L. Bartram, "Exploiting HMM Sparsity to Perform Online Real-Time Nonintrusive Load Monitoring," *IEEE Transactions on Smart Grid*, vol. 7, no. 6, pp. 2575–2585, Nov. 2016, <https://doi.org/10.1109/TSG.2015.2494592>.

[14] A. F. Ebrahim and O. A. Mohammed, "Energy Disaggregation Based Deep Learning Techniques: A Pre-processing Stage to Enhance The Household Load Forecasting," in *2018 IEEE Industry Applications Society Annual Meeting (IAS)*, Portland, OR, USA, Sep. 2018, <https://doi.org/10.1109/IAS.2018.8544664>.

[15] C. Zhang, M. Zhong, Z. Wang, N. Goddard, and C. Sutton, "Sequence-to-point learning with neural networks for nonintrusive load monitoring," *arXiv:1612.09106 [cs, stat]*, Sep. 2017, Accessed: May 07, 2021. [Online]. Available: <http://arxiv.org/abs/1612.09106>.

[16] M. Kaselimi, E. Protopapadakis, A. Voulodimos, N. Doulamis, and A. Doulamis, "Multi-Channel Recurrent Convolutional Neural Networks for Energy Disaggregation," *IEEE Access*, vol. 7, pp. 81047–81056, 2019, <https://doi.org/10.1109/ACCESS.2019.2923742>.

[17] R. Bonfigli, A. Felicetti, E. Principi, M. Fagiani, S. Squartini, and F. Piazza, "Denoising autoencoders for Non-Intrusive Load Monitoring: Improvements and comparative evaluation," *Energy and Buildings*, vol. 158, pp. 1461–1474, Jan. 2018, <https://doi.org/10.1016/j.enbuild.2017.11.054>.

[18] V. Piccialli and A. Sudoso, "Improving Non-Intrusive Load Disaggregation through an Attention-Based Deep Neural Network,"

- Energies*, vol. 14, no. 4, Feb. 2021, Art. no. 847, <https://doi.org/10.3390/en14040847>.
- [19] L. B. Salah and F. Fourati, "Systems Modeling Using Deep Elman Neural Network," *Engineering, Technology & Applied Science Research*, vol. 9, no. 2, pp. 3881–3886, Apr. 2019, <https://doi.org/10.48084/etasr.2455>.
- [20] S. Abid, M. Chtourou, and M. Djemel, "A Pruning Algorithm Based on Relevancy Index of Hidden Neurons Outputs," *Engineering, Technology & Applied Science Research*, vol. 6, no. 4, pp. 1067–1074, Aug. 2016, <https://doi.org/10.48084/etasr.704>.
- [21] J. Kelly and W. Knottenbelt, "Neural NILM: Deep Neural Networks Applied to Energy Disaggregation," in *Proceedings of the 2nd ACM International Conference on Embedded Systems for Energy-Efficient Built Environments*, New York, NY, USA, Nov. 2015, pp. 55–64, <https://doi.org/10.1145/2821650.2821672>.