Demand Side Flexibility Envelope Quantification Under Data Scarcity

Liudong Chen[1, 2], Xiangqi Zhu[2], Bolun Xu[1], and Fei Ding[2]
[1]Earth and Environmental Engineering, Columbia University, New York, USA, {lc3671, bx2177}@columbia.edu
[2]Power Systems Engineering Center, National Renewable Energy Laboratory, Golden, USA, {xiangqi.zhu, fei.ding}@nrel.gov

***Abstract*—Efficient management of distributed energy resources often requires a tremendous amount of data to model component constraints and consumer behaviors. However, in many practical scenarios, collecting high-resolution demand data is significant costly, which jeopardizes the efficient utilization of demand-side resources. This paper proposes a novel learning framework to quantify demand-side flexibility envelope under data scarcity. We construct a long-short-term memory network to capture the relationship between consumers' load consumption behavior and demand flexibility envelope. The network only uses low-resolution load data to predict consumers' personalized flexibility envelope and able to extrapolate across various stakeholders and conditions. We provide a flexibility envelope generation algorithm to generate the training label and expedites the training process with limited high-resolution historical load data and environmental data. Simulations using Pecan Street consumer data demonstrate that our method generate effective demand flexibility envelopes for network training, and the cross-validation results shows our learning framework accurately predict the demand flexibility envelope despite data scarcity.**

***Index Terms*—Demand response; data scarcity; quantification; neural network; flexibility envelope**

I. Introduction

Distributed energy resources (DERs) are the cornerstone of future power system operations. DERs offer consumers new options for electricity consumption and provide new market mechanisms to incentivize consumer engagements such as demand response (DR) programs enable [1], [2]. One of the key factors in designing an effective DR programs is modeling consumers' demand flexibility, represented by their upper and lower demand envelope [3], [4]. Non-intrusive load modeling (NILM), also known as load disaggregation, offers a convenient and privacy-ensuring approach to quantify the individual contribution of appliances to the load and to enable a detailed examination of consumers' demand patterns [5]. Some NILM research focuses on clustering and detecting load events, recognizing appliances' state transition processes as load signatures, and employing combinatorial optimization for load event matching [5], [6]. Other NILM methods utilize machine learning approaches to classify different types of appliances with supervised labels. The temporal and spatial features can also be captured by deep neural networks, which is widely used and proven efficient under the NILM framework. Transfer learning is also employed to enhance the network's extrapolation ability between different appliances [7]. Thus, with data support from NILM, utilities can potentially model consumers' demand flexibility based on detailed demand patterns [8].

However, NILM, while non-intrusive, often relies on high-resolution load data of minute or sub-minute sampling rate. While some low-resolution methods (15 minutes) necessitate extensive detailed datasets of historical appliance consumption [9]. These data sets require significant hardware investments to obtain high-frequency measurements via smart meters, and to store and process these data through cloud services. NILM algorithms require substantial data processing [4], leading to high energy consumption and not aligning with the principles of green computing [10]. Moreover, as NILM solely provides detailed appliance consumption patterns without in-depth consideration of human factors, it may fail to accurately reflect consumers' flexibility envelopes [11]. In practice, most DR programs operate on an hourly basis [12] and calculate demand flexibility using appliances' rated power or sending commands asking consumers to turn off certain appliances [13], [14]. Also, most of the utilities can only provide smart meter data up-to half-hour resolution, because they either don't have meters which can provide high resolution data, or data below half-hour cannot be disclosed due to privacy issue.

Therefore, predicting consumers' flexibility envelopes using low-resolution (half-hour or hourly-level) data not only efficiently utilizes existing DR communication and metering infrastructure, but also reduces the energy consumption and carbon footprint. Existing research about developing flexibility envelopes mainly focuses on the building sector, using performance indicators and flexibility functions to quantify the flexibility envelopes [15]. These model-based methods generally require specific modeling of each appliance [16]. Some works also use data-driven methods with detailed appliance consumption data [8] and historical data-driven quantification results [4]. Without heavily relying on the data, this paper designs a novel learning-based framework to quantify demand side flexibility envelopes under high-resolution data scarcity. The proposed framework leverages consumers' historical load consumption behaviors and enables utilities to rapidly predict flexibility envelopes using only hourly load measurements, circumventing the need for high-resolution data. The paper contributes in the following ways:

1) We propose a flexibility envelope learning framework that quantifies consumers' personalized demand flexibility en-

velopes by training a neural network with low-resolution load data. In this training process, demand flexibility envelopes serve as training labels, while historical load data and environmental data act as input features.
2) We develop a flexibility envelope generation algorithm that rapidly generates training labels and expedites the training process. This includes a fast load recognition algorithm to identify specific load patterns from high-resolution load profiles.
3) We use a real dataset to assess our learning framework, and the cross-validation results demonstrate the high accuracy of our envelope prediction outcomes.

The remaining of the paper is organized as follows: Section II introduces the flexibility envelope learning framework. Section III presents the flexibility envelope generation algorithm. Section IV introduce the dataset and computation results. Section V concludes the paper with a discussion on future directions.

II. Flexibility Envelope Learning Framework

A. Framework Overview

The overall framework is shown in Fig. 1. Our framework aims to predict demand flexibility envelope directly using hourly resolution data. Based on our discussions with three different utilities and two meter vendors across the US, it is a reasonable assumption that utilities have access to high-resolution demand measurement of limited consumers to train the prediction model. Hence, the trained model can be directly applied to a large number of consumers without prior high-resolution measurements. During the network training process, we use high-resolution data (minutes level) to generate the training label. We provide a fast load recognition method by separating consumers' load into three types: heating, ventilation, and air conditioning (HVAC)-driven load, behavioral-driven load, and base load, and we recognize each type's pattern separately. Among them, HVAC- and behavior-driven load are changeable and can provide flexibility. Based on that, we propose a flexibility envelope generation algorithm to calculate high-resolution demand flexibility envelopes and convert it to a low-resolution envelope as the training label.

B. Learning Model

Considering the time series characteristics of consumers' load and envelope, we opt for a long-short term memory (LSTM) network, which can capture the inherent temporal relationships in time series data. Since utility companies typically know consumers' hourly load consumption patterns and events that require demand side flexibility, such as demand response, to generate hourly signals, we build the learning model based on low-resolution load profiles:

$$\boldsymbol{Y} = f(\boldsymbol{X}) \tag{1}$$

$$\boldsymbol{X} = [\boldsymbol{l}, \boldsymbol{T}] \tag{2}$$

$$\boldsymbol{Y} = [E_h^+, E_h^-] \tag{3}$$

Where $\boldsymbol{Y}, \boldsymbol{X}$ are the output labels and input features; $\boldsymbol{l}$ is a matrix with hourly load consumption of historical days and target days; $\boldsymbol{T}$ is the target days' outdoor temperature; E_h^+, E_h^- are the upper and lower flexibility range.

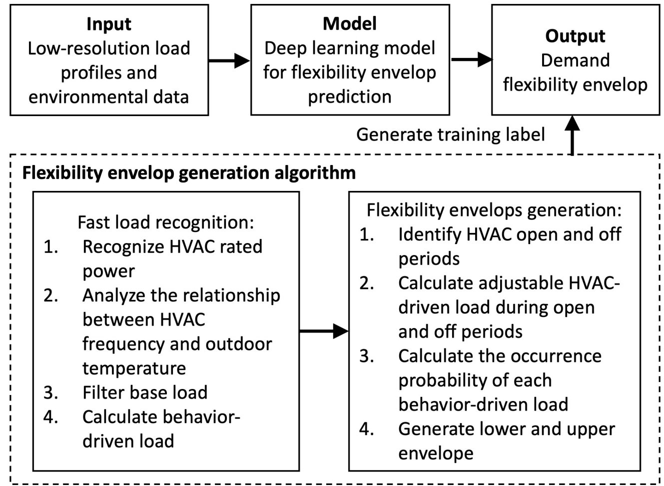

Fig. 1. Flexibility envelope learning framework.

To predict the demand flexibility envelope on target days, we use historical load profiles during several months and the load profiles of the specific target day to learn consumers' consumption behaviors and identify the potential for demand adjustments. Moreover, historical and target day outdoor temperatures play a dual role. The former assists in understanding the relationship between a user's consumption behavior and outdoor temperatures, while the latter contributes to determining the load profiles for the target day. It is noted that we only focus on predicting flexibility envelopes for target days, and the load profiles for these days can be obtained from various existing load forecasting methods.

We calculate the upper and lower envelope by adding or subtracting a flexibility range from the original load profiles. This range primarily comes from the adjustable portion of consumers' load. Also, the accuracy of the entire envelope prediction is affected by the precision of the original load prediction. Thus, here we predict the flexibility range E_h^+, E_h^-, rather than the entire envelope, i.e., the difference between the upper/lower envelope and the original load. We provide a flexibility envelope generation algorithm to calculate the flexibility range based on historical days' load profiles, serving as the training label, which we detail in the next Section.

With consumers' flexibility ranges determined, we merge consumers' original load profiles on the target days with the upper flexibility range to calculate the entire upper envelope. This approach considers the presently active appliances, noting that appliances already in operation cannot contribute to the current upper bound. For the entire lower envelope, we add the lower flexibility envelope to the base load to obtain the minimum fixed load.

The LSTM network learns the relationship between historical load and flexible envelope by recognizing the load patterns and flexibility range from the input features. Once we learn enough consumers' behavior from the labels calculated using high-resolution load data, the network can extrapolate to other consumers with only low-resolution load data. The patterns and trends of consumers' load profiles are similar, which support the network extrapolation. However, note that social demographics significantly influence consumers' behavior. In this context, we assume consumers with similar load profiles will provide similar load flexibility. With more social demographics data, we can include this information in the network training process to better capture consumers' individual flexibility. Also, the trained network can be updated with more recent load profiles by periodically training new networks or using the transfer learning method.

III. Flexibility Envelope Generation Algorithm

A key component in our training process is to generate flexibility envelopes from historical high-resolution demand measurements, which are used as training labels. To this end, we introduce a method to fast recognize load patterns and develop an algorithm to calculate the demand side flexibility envelope as the training label of the learning model.

A. Fast Load Recognition

We separate load into three types as shown in Fig. 2. Among them, an HVAC-driven load is a special large adjustable load, usually ranging from 1-4kW, regularly appearing throughout the day. HVAC operation is influenced by the outside temperature and consumers' indoor temperature set points. When indoor temperatures reach the set points, the HVAC will turn off; otherwise, it remains operational. Thus, rated power and frequency are commonly used to express the HVAC-driven load.

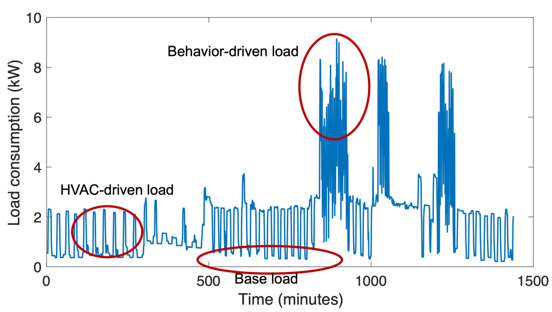

Fig. 2. Load patterns from high-resolution load profiles.

Behavior-driven load is determined by consumers' uncertain behavior, which is irregular with fast variations, such as hairdryers, washing machines, or cooking appliances. In general, most behavior-driven loads can be adjusted to provide flexibility. On the other hand, the base load encompasses the standby energy consumption of appliances, characterized by its stability and relatively small magnitude. This load does not offer flexibility during incentive or mandatory events. Consequently, we employ a minimal smoothing method to filter the base load:

$$l_{\mathrm{base},t'} = \min(l_t, ..., l_{t+t_w}),\ \forall t' \in [t, t+t_w] \tag{4}$$

where l_t is the entire load consumption at time slot t; $l_{\mathrm{base},t'}$ is the base load in time slot t', t_w is the window size in the smoothing method.

Learned from the specific patterns exhibited by the HVAC-driven load, which typically manifests as a spike characterized by a substantial increase over 2-4 minutes, followed by a sustained plateau during subsequent brief intervals dictated by the HVAC frequency, we first provide the following criteria to recognize HVAC rated power l_{HVAC}.

$$l_{t+1} - l_t \geq \tau_1 \tag{5}$$

$$\min(l_{t+2}, l_{t+3}, l_{t+4}) - l_t \geq \tau_2 \tag{6}$$

$$|\min(l_{t+5}, ..., l_{t+t_{\mathrm{s}}}) - \max(l_t, ..., l_{t+4})| \leq \tau_3 \tag{7}$$

where t_{s} is the time span for the HVAC spike, τ_1, τ_2, τ_3 are the thresholds used to judge the spike in consumers' demand, which is determined by the demand profiles.

Equation (5) recognizes the beginning of a spike, (6) judges the increase stage of a spike, and (7) judge if the spike stays for a short time, with the time span determined by the HVAC frequency. However, the behavior-driven load can also exhibit similar patterns due to its highly random and personal nature. Since the behavior-driven load is irregular, we use a histogram to calculate the frequency of each spike appearing and choose the most frequent spike value as the HVAC rated power. Then we count the frequency of the HVAC-driven load, expressed as the percentage of HVAC on-time during each hour, and use linear regression to establish its relationship with the outside temperature:

$$f_{\mathrm{HVAC},T} = \mathbb{E}_{\boldsymbol{f}_{\mathrm{HVAC}} \in T}[\boldsymbol{f}_{\mathrm{HVAC}}] \tag{8}$$

$$f_{\mathrm{HVAC},T} = a * T + b \tag{9}$$

where a, b are the linear regression parameters; T is the outdoor temperature; $\boldsymbol{f}_{\mathrm{HVAC}}$ is the HVAC frequency over time, and $f_{\mathrm{HVAC},T}$ is the average HVAC frequency corresponding to outdoor temperature.

After obtaining the HVAC-driven load and base load, we subtract them from the consumer's entire load to obtain the behavior-driven load.

B. Flexibility Envelope Generation

Based on the recognized HVAC-driven load, behavior-driven load, and base load patterns, we can formulate consumers' demand flexibility envelope. This flexibility is expressed by the upper and lower envelope of consumers' demand, indicating the demand change capacity, contributed by the HVAC-driven load and behavior-driven load.

Regarding the HVAC-driven load, consumers can adjust the indoor temperature set point to adjust the on/off time periods to provide demand flexibility. For example, in cooling zone, during the on time periods, consumers can increase the indoor temperature set point to reduce load consumption, which contributes to the lower envelope. Conversely, during the off time periods, consumers can reduce the set point to open the HVAC and consume energy, providing the upper envelope. The detailed formulation is expressed as follows:

$$E^{+}_{\mathrm{HVAC},t} = s_t * l_{\mathrm{HVAC}},\ \forall t \in t_{\mathrm{off}} \tag{10}$$

$$E^{-}_{\mathrm{HVAC},t} = (1 - s_t) * l_{\mathrm{HVAC}},\ \forall t \in t_{\mathrm{on}} \tag{11}$$

where s_t is the allowed set point change by consumers in time slots t, for convenience, expressed as a percentage of consumers' energy change; $t_{\text{off}}, t_{\text{on}}$ are the HVAC off and open time slots, respectively; $E^{+}_{\text{HVAC},t}, E^{-}_{\text{HVAC},t}$ are the upper and lower envelope contributed by the HVAC-driven load.

The behavior-driven load consists of many small appliances with uncertain consumer behaviors, making it challenging to recognize granular load patterns under data limitation. To address this issue, we provide a statistical method based on historical load profiles to quantify the contribution of behavior-driven load in the upper and lower envelope. We obtain historical behavior-driven load values for each minute of each day using high-resolution load profiles. This allows us to calculate the occurrence probability of specific behavior-driven loads. We compute the expected value of historical behavior-driven load that surpasses the behavior-driven load for target days. This calculation indicates the likelihood of higher loads occurring during those specific time slots. Conversely, for the lower envelope, we calculate the minimum behavior-driven loads from both historical and target days, which should be treated as fixed loads as they always exist. The equation is then expressed as follows:

$$E^{+}_{\text{B},t} = \mathbb{E}[\boldsymbol{l}_{\text{B},t} - l_{\text{Bt},t}],\ \ E^{-}_{\text{B},t} = \min([\boldsymbol{l}_{\text{B},t}, l_{\text{Bt},t}]), \qquad (12)$$

where $\boldsymbol{l}_{\text{B},t}$ is the behavior-driven load matrix in historical days at minute t; $l_{\text{Bt},t}$ is the behavior-driven load vector in the target days; $[*]$ means the concatenation of matrix.

We then calculate the flexibility range determined by the HVAC- and behavior-driven loads across historical and target days. We first convert the high-resolution load profiles into a low-resolution representation. The HVAC-driven load is calculated by the hourly outdoor temperature and rated power because the variations in indoor and outdoor temperature differences affect the HVAC open time in each time slot, which can be expressed as follows:

$$l_{\text{V},h} = l_{\text{HVAC}} * (a * T_h + b) \qquad (13)$$

where $l_{\text{V},h}$ is the HVAC-driven load at hour h.

The hourly flexibility range, serving as the training label, is expressed as follows:

$$E^{+}_{h} = s_h * (l_{\text{HVAC}} - l_{\text{V},h}) + \mathbb{E}[\boldsymbol{l}_{\text{B},h} - l_{\text{Bt},h}] \qquad (14)$$

$$E^{-}_{h} = (1 - s_h) * l_{\text{V},h} + \min([\boldsymbol{l}'_{\text{B},h}, l_{\text{Bt},h}]) \qquad (15)$$

IV. Case Study

A. Datasets

In the case study, we used a dataset from Pecan Street. This dataset contains one year of minute-level load consumption data for 148 consumers in Texas, USA. Additionally, hourly temperature data for the same area is collected from the National Solar Radiation Database of the National Renewable Energy Laboratory (NREL) [17]. To increase the number of consumers for better network training, we separated the yearly data into monthly data and assumed that each month represents a different consumer so that we increased the number of consumers to 1776. To implement our algorithm on these

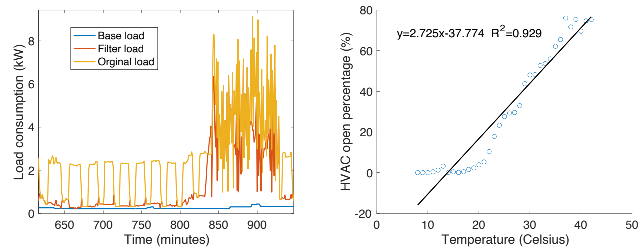

(a) Three load patterns recognition (b) HVAC frequency with temperature

Fig. 3. Load recognition results.

consumers, we specified the following parameters: τ_1, τ_2, τ_3 defined in equation (5), (6), and (7) above as 0.2kW, 1.5kW, 0.4kW,, t_w, s_t, t_{s} defined in equation (4), (10), and (7) above as 180 minutes, 80%, and 6 minutes, respectively. Note that the first three threshold parameters and the HVAC spike time span are recognized from the load profiles, and set point parameters can be accurately learned if consumers' set point information is included.

B. Flexibility Envelope Generation Results

We first demonstrate the effectiveness of our load recognition algorithm used to generate the learning training label. Fig. 3 (a) shows the load recognition results given the original high-resolution load profiles. The filtered load effectively excludes the HVAC-driven load, as every regular spike is recognized and reduced from the original load. However, due to the possibility of overlap between behavior-driven and HVAC-driven loads, some loads with HVAC-rated power are also reduced, even if they don't exhibit regular patterns in the entire load profiles. For instance, in time slots 850-900, the HVAC-driven load patterns are covered by the behavior-driven load. The relationship between outdoor temperature and HVAC open time percentage is depicted in Fig. 3 (b), showing a strong linear relationship with an R^2 value of 0.93.

Fig. 4 shows the load flexibility envelope results from the load recognition results. We first compress the high-resolution load profiles into a low-resolution level to develop practical and useful low-resolution envelope. From the envelope, it is evident that during the original peak time slots, the difference between the upper envelope and the original load is small, while the difference between the lower envelope and the original load is large. This is a reasonable result, as small loads can be shifted to these peak time slots since consumers already use many appliances, while larger loads can be shifted to other time slots. Additionally, the maximum value of the upper envelope lies within the consumers' whole rated power, as it is calculated by the average value of the historical load. The lower envelope also always keeps a small load range according to consumers' load consumption behavior, providing for the basic energy requirement of consumers.

C. Flexibility Envelope Prediction Results

This section cross-validates the LSTM prediction performance for the upper and lower flexibility envelope. We use 29 days of low-resolution load profiles and hourly temperature

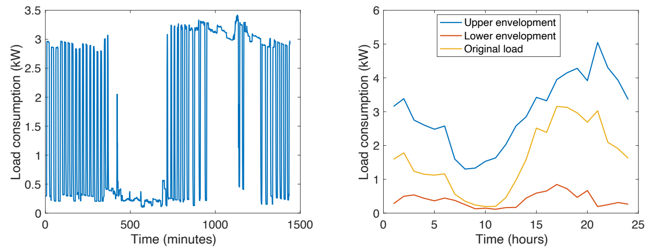

(a) High-resolution original load profiles

(b) Low-resolution upper and lower envelope

Fig. 4. Upper and lower envelope development results

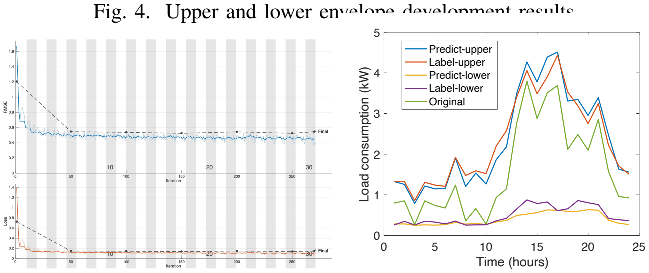

(a) Learning process

(b) envelope prediction

Fig. 5. Upper and lower envelope prediction results.

data from the 29th day as the input, while the 29th day's flexible upper and lower envelope served as the output to train the LSTM network. The network is constructed with 30 input features and 24 time steps. It includes a sequence input layer, an LSTM layer, a fully connected layer, and a regression layer. We set the number of neurons to 150 and the minimal batch size to 64. We chose Adam as the gradient-based optimizer and use MATLAB 2022b to build the network. The algorithm is implemented on a MacBookPro with an Apple M2 chip, 16 GB memory, and the training time is about 20s.

Consumers from the dataset are divided into 70%, 10%, and 20% for training, validation, and testing, respectively. Fig. 5 (a) shows the training process of the LSTM model, demonstrating good convergence and accuracy. In addition to using the built-in validation of the LSTM toolbox to enable early stopping, we also validate the network using k-fold cross-validation with $k = 5$, the validation RMSE reaches 0.22, the MAE is 0.17, and the loss is 0.07. The envelope prediction results in the test set are shown in Fig. 5 (b), highlighting the model's ability to accurately capture the correlation between demand flexibility and historical load consumption, even with limited data (low-resolution data). Although minor fluctuations are observed, demand side flexibility envelopes are generally used after clustering consumers to provide grid services or join the market to make profits. These fluctuations can be offset during consumer clustering, providing stakeholders with accurate demand side flexibility envelope information.

V. Conclusion

This paper proposes a novel learning framework to quantify demand side flexibility with limited data, applicable to various conditions that require consumers' demand adjustments, e.g., DR, resilience enhancement, system planning, and market design. Using high-resolution data from Pecan Street, we recognize the patterns of HVAC-driven load, behavior-driven load, and base load, and use them to generate consumers' demand flexibility envelope. The envelope aligns reasonably with historical consumption behaviors and demand shift possibilities. Furthermore, the flexibility envelope prediction results demonstrate the effectiveness of our learning framework in recognizing the envelope composition from low-resolution historical load profiles, achieving a RMSE of 0.22 by cross-validation. Future research will focus on exploring advanced machine learning methods to accurately recognize flexibility, considering more consumers' behavior and social demographics, such as comfortable range, indoor temperature set point, income, and age, to develop better flexibility envelope. Additionally, as this work assesses how much consumers can respond, another direction is to investigate consumers' willingness to respond and the incentives they require for demand adjustments.

References

[1] I. Perez-arriaga et al, "Utility of the future: An mit energy initiative response to an industry in transition," Massachusetts Institute of Technologye, Tech. Rep., 2016.

[2] N. Li, L. Chen, and S. H. Low, "Optimal demand response based on utility maximization in power networks," in *2011 IEEE Power and Energy Society General Meeting*, 2011, pp. 1–8.

[3] Lawrence Berkeley National Laboratory, "Demand flexibility." [Online]. Available: https://buildings.lbl.gov/demand-flexibility

[4] N. Hekmat, H. Cai, T. Zufferey, G. Hug, and P. Heer, "Data-driven demand-side flexibility quantification: Prediction and approximation of flexibility envelopes," *arXiv preprint arXiv:2110.12796*, 2021.

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

[6] B. Liu, W. Luan, J. Yang, and Y. Yu, "The balanced window-based load event optimal matching for nilm," *IEEE Transactions on Smart Grid*, vol. 13, no. 6, pp. 4690–4703, 2022.

[7] M. D'Incecco, S. Squartini, and M. Zhong, "Transfer learning for non-intrusive load monitoring," *IEEE Transactions on Smart Grid*, vol. 11, no. 2, pp. 1419–1429, 2019.

[8] H. Cai and P. Heer, "Experimental implementation of an emission-aware prosumer with online flexibility quantification and provision," *arXiv preprint arXiv:2110.12831*, 2021.

[9] B. Zhao, M. Ye, L. Stankovic, and V. Stankovic, "Non-intrusive load disaggregation solutions for very low-rate smart meter data," *Applied Energy*, vol. 268, p. 114949, 2020.

[10] P. Kurp, "Green computing," *Communications of the ACM*, vol. 51, no. 10, pp. 11–13, 2008.

[11] K. C. Armel, A. Gupta, G. Shrimali, and A. Albert, "Is disaggregation the holy grail of energy efficiency? the case of electricity," *Energy policy*, vol. 52, pp. 213–234, 2013.

[12] U.S. Department of energy, "Demand response." [Online]. Available: https://www.energy.gov/oe/demand-response

[13] Tesla, "Tesla demand response terms of service." [Online]. Available: https://www.tesla.com/support/energy/more/legal/demand-response-terms-of-service

[14] New York Independent System Operator, "Demand response." [Online]. Available: https://www.nyiso.com/demand-response

[15] R. Li, A. J. Satchwell, D. Finn, T. H. Christensen, M. Kummert, J. Le Dréau, R. A. Lopes, H. Madsen, J. Salom, G. Henze *et al.*, "Ten questions concerning energy flexibility in buildings," *Building and Environment*, vol. 223, p. 109461, 2022.

[16] P. Munankarmi, X. Jin, F. Ding, and C. Zhao, "Quantification of load flexibility in residential buildings using home energy management systems," in *2020 American Control Conference (ACC)*, 2020, pp. 1311–1316.

[17] National Renewable Energy Laboratory, "National solar radiation database." [Online]. Available: https://nsrdb.nrel.gov/