Automated Classification of Appliances Using Elliptical Fourier Descriptors

Leen De Baets, Chris Develder, Tom Dhaene and Dirk Deschrijver
Department of Information Technology, Ghent University - imec
Technologiepark-Zwijnaarde 15, 9052 Ghent, Belgium
Email: leen.debaets@ugent.be

Abstract—**Non-intrusive load monitoring methods aim to dis-aggregate the total power consumption of a household into individual appliances by analysing changes in the voltage and current measured at the grid connection point of the household. The goal is to identify the active appliances, based on their unique fingerprint. This information can be communicated to the electricity provider and the end-user enabling the potential of smart grids. An informative characteristic to attain the appliance classification is the voltage-current trajectory. In this paper, this trajectory is represented as a binary image from which the contours are extracted. From these contours, the elliptic Fourier descriptors are calculated and used as input for several classification algorithms outputting the appliance name. Benchmarking this method on the PLAID dataset shows that the descriptors can yield a prediction accuracy up to 79%, comparable to the state-of-the-art, based on only a very compact representation (12 numbers).**

I. INTRODUCTION

Nowadays, the use of smart grids is becoming predominant [1], [2]. Some factors that motivate the transition to smart grids are, among others, (1) meeting renewable energy targets, (2) system efficiency improvements, (3) reliability improvements, (4) enabling customers choice and participation, (5) enabling new services, products and markets, and (6) improving energy efficiency [3]. To achieve these goals in terms of technology priorities, it is critical to develop renewable energy sources (RES), and means to integrate them in a way that deals with their uncontrollable variability. To cope with abundant RES, it is crucial to have demand response (DR) capabilities to shift power consumption in time, e.g., to periods with high availability of RES power. On a residential level, DR can be achieved by shedding high energy consuming appliances at the request of the provider. To this end, non-intrusive load moni-toring (NILM) is proposed to check whether these appliances are active and responsive to the shedding instructions [4], [5], [6]. NILM will detect status changes of household appliances (e.g., ON/OFF events) by analysing the total household energy consumption measured with a smart meter. This smart meter is located at the users' home so that the measurements are done local. Once the changes are detected (e.g., using robust event detection methods [7]), the measured signals can be used to recognize and classify the corresponding appliances. Using NILM, it is possible to communicate (1) to the energy provider which appliances are active and responsive to the shedding commands, and (2) to the end user the appliance usage statistics.
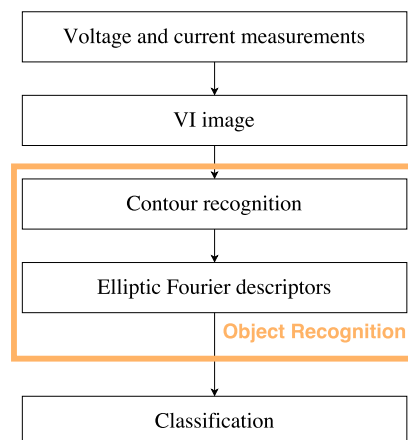

Fig. 1: The classification workflow.

In this paper, the problem of classifying appliances based on their voltage-current trajectories is addressed by apply-ing object recognition algorithms on the binary image of these trajectories. The next section covers the related work concerning these images and appliance recognition. Section III explains the different steps on how to describe objects in images: (1) find the contours, (2) calculate the elliptic Fourier descriptors of the contours. Section IV describes three machine learning methods used for classifying the appliances using the elliptic Fourier descriptors. Section V benchmarks the suggested method on the PLAID dataset and Section VI concludes the paper. The approach of this paper is visualised in Figure 1.

II. RELATED WORK

Classifying active appliances is mostly done by extracting features from the monitored data and training a machine learning classifier. The type of features depends strongly on the sampling rate of the measurements. When using low frequency data ($\leqslant 1$ Hz), the most common features are the power levels and the ON/OFF durations [8]. A drawback of this approach is that only energy-intensive appliances can be detected. This can be alleviated by performing fine-grained measurements at the cost of an increased data storage rate and more complex data analytics. It is then possible to calculate features such as the harmonics [9] and frequency components [10] from the steady-state and transient behavior of the current and voltage signal.
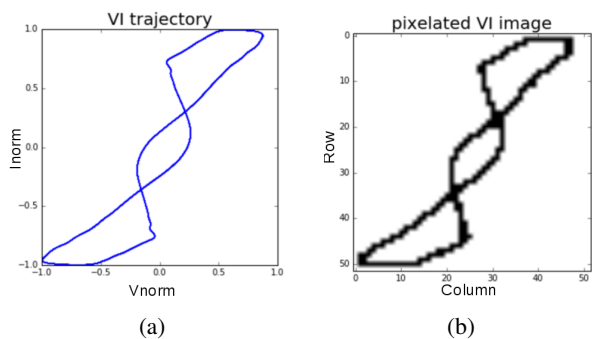
Fig. 2: The original VI trajectory (left) and the corresponding binary VI image (right).

Fig. 3: The lookup table consisting of sixteen different cases in which parts of the contour is formed.

More recently, the possibility to consider voltage-current (VI) trajectories has also been considered [11], [12], [13], [14].

The VI trajectory of an appliance is obtained by plotting the voltage against the current for a defined time period when the appliance is turned on, see Figure 2a. It is shown in [11] that manually extracting features from the VI trajectory such as the enclosed area, slope of the middle segment, etc. can be informative to classify the appliances.

Nevertheless, extracting features from the VI trajectory is not straightforward. As an alternative, the VI trajectory can be converted into a binary VI image ($n \times n$ matrix) by meshing the VI trajectory, see Figure 2b. In [12], [13], each cell of the mesh is assigned a binary value that denotes whether or not it is traversed by the trajectory. Based on this binary VI image, several features can be extracted to classify different power loads [13]. Examples of features are the number of continuums of occupied cells, or the binary value of the left horizontal cell and central cell. In [12], the binary VI image is transformed into an input vector that can be fed directly into a classifier like random forests to classify different appliances.

In order to distinguish appliances based on their VI trajectories, powerful measurement devices must be used that are able to sample high frequency data.

In this paper, object recognition methods are used to recognize the contours of a VI trajectory in the binary image, and to describe them using elliptical Fourier descriptors, see Section III.

III. OBJECT RECOGNITION

Object recognition is a fundamental problem in computer vision and its concepts have been applied in multiple fields, such as distinguishing biological species [15], optical character recognition [16], and face recognition [17]. In object recognition, the contours of an object are identified from the image, characterised by elliptic Fourier descriptors and then classified with a label. In this paper, object recognition methods are used to recognize the contours of a VI trajectory in the binary image, and to describe them using elliptical Fourier descriptors. The next subsections explain these steps.
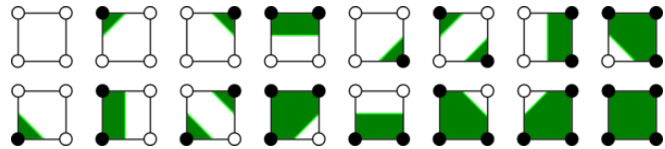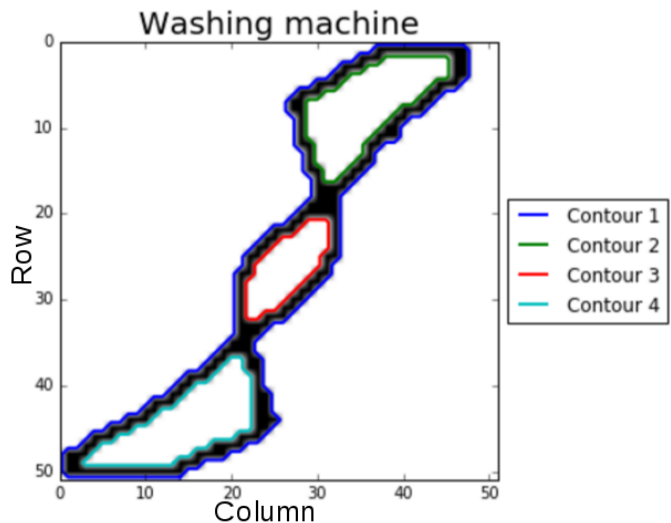

Fig. 4: Identification of the contours from VI trajectory of a washing machine. Only the outer contour is used for object classification.

A. Contours

The contour of an object in an image is a closed curve that forms the boundary of that object. The marching squares algorithm [18] can be used to find all contours in an image. The basic idea is that every pixel of the image is examined and matched with one of the sixteen cases shown in Figure 3. Each case creates at most two edges. Similar results can be obtained using other contour algorithms [19].

Figure 4 shows the detected contours of the VI trajectory of a washing machine. This example has four contours, from which only the outer one is of interest. To avoid that the trajectory touches the border, extra pixel rows and columns are added to the sides. Otherwise this would result in two separate outside contours instead of one.

The outer contour is a closed curve that resembles the shape of a smoothed VI trajectory. In contrast to the original VI trajectory, all points on the contour are separated uniformly, such that the Euclidean distance between neighbouring points on the contour is the same.

B. Elliptic Fourier descriptors

Once the contour is identified, elliptical Fourier descriptors (EFD) are used to match it with the corresponding appliance. A brief outline is given in this section, and the reader is referred to [20], [21] for mathematical details.
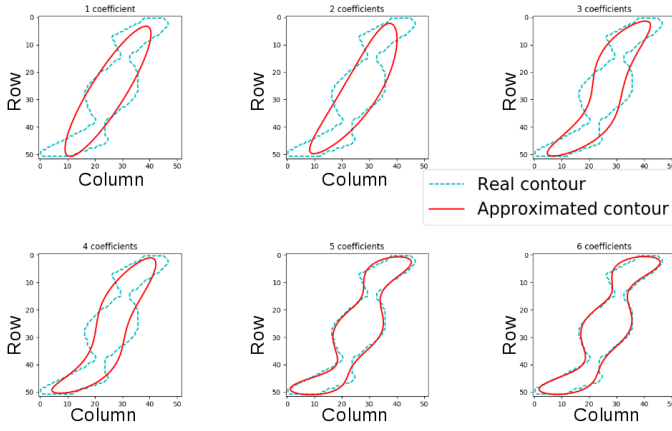
Fig. 5: Original (blue) and approximated (red) contours of the VI trajectory of an appliance with increasing number of coefficients.

EFDs define the contour as the sum of a certain number of ellipses (E) required to mimic the shape. Each ellipse is formed by two sets of partial differential equations, each having sine and cosine terms. One set of equations is defined along the x-axis and the other one along the y-axis. This is a transformation from the spatial domain to the frequency domain:

$$x(t) = A_0 + \sum_{e=1}^{E-1} A_e\cos(et) + B_0 + \sum_{e=1}^{E-1} B_e\sin(et) \quad (1)$$

$$y(t) = C_0 + \sum_{e=1}^{E-1} C_e\cos(et) + D_0 + \sum_{e=1}^{E-1} D_e\sin(et) \quad (2)$$

where A_0, B_0, C_0, and D_0 are the constants, A_e, B_e, C_e, and D_e are the harmonic coefficients of the e^{th} order, and t represents the points sampled from the time axis given by the period 2π. B_0 and D_0 are zero and can be omitted from the formula.

Each harmonic is thus described by four Fourier coefficients, two each for the x- and y-axis, generating a total of $4 \cdot E$ coefficients. The first harmonic describes the overall shape, location, size, and rotational orientation of the contour. Additionally, consecutive harmonics can be included to capture more detailed information about the contour's complexity. Figure 5 shows the reconstruction of the contour when using up to $E = 6$ harmonics. The approximated contour better resembles the original contour when more harmonics are included. Using $4 \cdot E$ EFDs, it is possible to recreate the contour without requiring prior expertise on the important features to represent it.

IV. OBJECT CLASSIFICATION

The object recognition results in a vector of size $4 \cdot E$. This vector can be used as input for classification algorithms. This paper focuses on three methods: logistic regression, random forests and a simple neural network (non-linear model).

A. Logistic Regression

For a binary problem, logistic regression (LR) predicts the probability that a sample belongs the '1' class versus the probability it belongs to the '0' class given one or more independent input variables. For the multi-class problem, a binary logistic regression is fitted for each label, creating a one-versus-all solution. The error function for a binary problem with m samples is:

$$Cost = -\frac{C}{m}\sum_{i=1}^{m} y_i \log(\hat{y}_i) + (1-y_i)\log(1-\hat{y}_i) + \sum_{i=1}^{k} w_i^2 \quad (3)$$

where y_i and $\hat{y}_i$ are respectively the real and predicted output of sample i, and C is a trained regularization parameter that controls the model complexity. When C is small, the first sum in the error function is smaller implying that the learning is slower and the regularization is stronger. The second sum consists of the trainable squared weights w_i, one for each input feature. This sum forms an additional method to prevent overfitting, called L2 regularization, implying that a solution with small weights is preferred over a solution with large weights.

B. Random Forest

A random forest (RF) is an ensemble technique that classifies the data using a collection of decision trees. Each decision tree is trained on a subset of the dataset that has the same size as the original training set, but samples are drawn with replacement. At each node of the decision tree, a feature is selected and the tree is traversed downwardly (either following left/right branch) by comparing its value to a threshold. Given a new sample, the output of each decision tree is averaged to obtain the final prediction.

The parameters that must be trained, are the amount of trees t and the maximal amount of features f that are considered in each node to decide upon what the optimal split is. In our case, f is at most $4 \cdot E$, the total amount of coefficients representing the contour.

C. Neural Networks

The architecture of a neural network (NN) consists of different layers, see Figure 6. The first layer is the input layer containing as many nodes as the dimension of a sample (here, $4 \cdot E$). This is followed by one (or more) fully connected layers which are hidden. Each of these layers contains a certain amount of nodes that have learnable weights and biases, which define a linear transformation of the node's inputs, after which a non-linearity is performed optionally. This non-linearity is often obtained by using a rectified linear unit that replaces all negative values by zero. At the end, the output of the last fully connected layer is fed into the output layer. Since the NN is used for classification, the output layer has K nodes with K equal to the amount of classes. The values of the output nodes are chosen to lie between 0 and 1 and sum up to 1, which is achieved by applying the softmax function. In other words, each node represents the probability that the EFDs corresponds
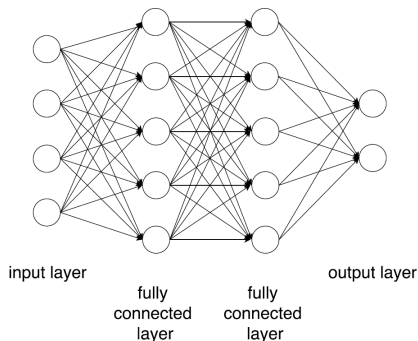
Fig. 6: A neural network consisting of an input layer, two fully connected layers and an output layer.

Method	Parameter	Range
LR	C	$[10^9, 10^7, 10^5, 10^3, 10^2, 1, 0.01]$
NN	α	$[10^{-5}, 10^{-3}, 1, 10]$
	$[h_1, ..., h_n]$	$[(10), (50), (100),$
		$(10, 10), (50, 50), (100, 100)]$
RF	t	$[10, 20, 30, 50, 100, 200]$
	f	$[4 : 4 : 4 * \text{amountCoefficients}]$

TABLE I: The considered parameter ranges.

to a given label. The output node with the maximal value represents the predicted class.

By providing training data, the weights of the network can be learned in such way that the predicted output is close to the real output. The learning rate α controls the speed of learning and is a parameter that needs to be trained. A smaller value reduces the risk of overfitting, at the cost of slower learning.

The parameters that are trained for neural network are the learning rate α, the amount of dense hidden layers and the amount of nodes in each layer. The last two parameters are symbolized by $(h_1, h_2, ..., h_n)$, where h_i represents the number of nodes in the hidden layer i, and n is the number of hidden layers.

V. RESULTS

A. Benchmark dataset

Ideally, to test the proposed method, a dataset having high frequency aggregated and high frequency sub-metered V and I signals would be used. However, to the best of our knowledge, no public dataset includes both. For this reason the PLAID [22] dataset is considered. Plug-Load Appliance Identification Dataset (PLAID) [22] is a public dataset including sub-metered current and voltage measurements sampled at 30 kHz for 11 different appliance types. More than 200 individual appliances are available, captured in 55 households. For each appliance, at least 5 start-up events are measured, resulting in a total of 1074 measurements.

The presented research on appliance classification is a first step towards a more realistic NILM setting starting from the aggregated power measurements. Still, this is a very meaningful one, as typically appliances are turned on/off one at a time, and the single appliance current (and thus VI trajectory) can be extracted from the aggregated measurements by considering the difference in current before and after the event. Future work needs to confirm the practical feasibility of this idea.

B. Calculation of EFD from VI images

In order to obtain the binary VI images, the voltage and current are collected over a time interval of 0.33 seconds, resulting in 10000 samples. In these images, the outer contour is detected from which the EFDs are calculated using a varying

number of harmonics. These components are used as input for three classification algorithms: logistic regression, random forests and neural networks.

C. Classification result

The generalization properties of the model are validated using leave-one-house-out cross-validation as recommended in [12]. Each training set contains data from 54 houses and the test set data from the remaining house. This is repeated 55 times. Using leave-one-house-out cross-validation allows to validate the generalization between different appliances of the same type as different houses, e.g., contain air conditioners but from different brands.

In order to train the parameters of the machine learning algorithms, grid search is performed. The ranges of the trained parameters are shown in Table I. In order to avoid overfitting while training these parameters, 10-fold cross-validation is performed. The 54 houses are partitioned into 10 equally sized subsets. Each time, 9 of those subsets are used for training and the remaining one is used for model validation. This process is repeated 10 times, with each of the 10 subsets used exactly once as the validation set. The 10 results can then be averaged to obtain a single prediction. The parameters resulting in the best validation score are selected and used on the single test house that was left out in the beginning.

In the end, there will be 55 test scores (one for each house). To obtain the final test result, these values are averaged. To evaluate the results, the following accuracy metric is used:

$$\text{Accuracy} = \frac{\#\text{correct predictions}}{\#\text{total predictions}} \quad (4)$$

Figure 7 shows the result when using the different machine learning methods and a varying amount of EFDs. The EFDs are calculated from contours that were extracted from 16×16 sized binary VI images. Using three or more EFDs does not significanty impact the accuracy for logistic regression and random forests. In the case of neural networks, the accuracy drops when more than 4 EFDs are used. The results of a random forest are comparable to those reported in [22], where the 16×16 binary VI images were used as plain input for a random forest (81.75% accuracy).

Figure 8 shows the confusion matrices when classifying the appliances using logistic regression, random forests and neural networks using 3 EFDs. The values in the matrix represent the
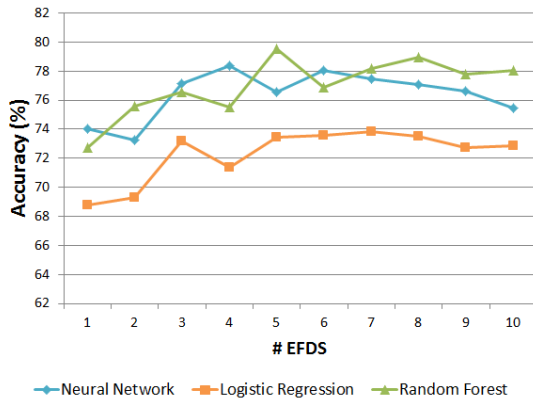
Fig. 7: Accuracy of the three classifiers

absolute number of appliances detected. The color represents per appliance the relative amount with respect to the total number of that appliance. For the three methods, appliance classes that are difficult to classify are the same. However, for random forest and neural network, the amount of correctly classified appliances is higher. The low performing classes are the heater, fridge, and AC. The heater is constantly confused with the hairdryer, explained by the fact that both contain a heating element. The fridge and AC are sometimes confused with each other and also with the fan, hairdryer and ILB.

A key advantage of this approach over [22] is that the required storage of the features is far less, albeit at a small additional cost of computing the contours. Only 12 values per sample should be stored compared to 256 values that are needed to store the original 16×16 binary image.

VI. Conclusion

In this paper, the contours of the binary VI image are used to characterise appliances in NILM. From these contours, the elliptic Fourier descriptors are calculated and used as input for logistic regression, random forests, and neural networks. The results show that twelve components per sample (i.e., three times four harmonic coefficients) are sufficient to obtain a prediction accuracy of 77.14%. This leads to significant ($> 20\times$) storage savings when compared to the original VI image comprising $16 \times 16 = 256$ values.

References

[1] G. Di Leo, C. Liguori, V. Paciello, A. Pietrosanto, and P. Sommella, "Smart meters and demand response: A double challenge," in *Computational Intelligence and Virtual Environments for Measurement Systems and Applications (CIVEMSA), 2015 IEEE International Conference on.* IEEE, 2015, pp. 1–6.

[2] C.-W. Tsai, A. Pelov, M.-C. Chiang, C.-S. Yang, and T.-P. Hong, "Computational awareness for smart grid: a review," *International journal of machine learning and cybernetics*, vol. 5, no. 1, pp. 151–163, 2014.

[3] W. P. Wang and E. Lightner, "International assessment of smart grid drivers and technologies: Analysis across countries and by economy and continent," in *PES General Meeting— Conference & Exposition, 2014 IEEE.* IEEE, 2014, pp. 1–5.

[4] K. Basu, V. Debusschere, and S. Bacha, "Residential appliance identification and future usage prediction from smart meter," in *Industrial Electronics Society, IECON 2013-39th Annual Conference of the IEEE.* IEEE, 2013, pp. 4994–4999.

(a)

(b)

(c)

Fig. 8: The confusion matrices when classifying the appliances using (a) logistic regression, (b) random forest, and (c) neural networks. CFL = compact fluorescent lamp, AC = air conditioner, ILB = incandescent light bulb

[5] D. L. Racines and J. E. Candelo, "Non intrusive load identification with power and impedance obtained from smart meters," *Int. J. Eng. Technol*, vol. 6, no. 4, pp. 1867–1876, 2014.

[6] I. Abubakar, S. Khalid, M. Mustafa, H. Shareef, and M. Mustapha, "An overview of non-intrusive load monitoring methodologies," in *Energy Conversion (CENCON), 2015 IEEE Conference on.* IEEE, 2015, pp. 54–59.

[7] L. De Baets, J. Ruyssinck, C. Develder, T. Dhaene, and D. Deschrijver, "On the bayesian optimization and robustness of event detection methods in nilm," *Energy and Buildings*, vol. 145, pp. 57–66, 2017.

[8] N. Henao, K. Agbossou, S. Kelouwani, Y. Dubé, and M. Fournier, "Approach in nonintrusive type i load monitoring using subtractive clustering," *IEEE Transactions on Smart Grid*, 2015.

[9] W. Wichakool, Z. Remscrim, U. A. Orji, and S. B. Leeb, "Smart metering of variable power loads," *IEEE Transactions on Smart Grid*, vol. 6, no. 1, pp. 189–198, 2015.

[10] H.-H. Chang, K.-L. Chen, Y.-P. Tsai, and W.-J. Lee, "A new measurement method for power signatures of nonintrusive demand monitoring and load identification," *IEEE Transactions on Industry Applications*, vol. 48, no. 2, pp. 764–771, 2012.

[11] T. Hassan, F. Javed, and N. Arshad, "An empirical investigation of vi trajectory based load signatures for non-intrusive load monitoring," *IEEE Transactions on Smart Grid*, vol. 5, no. 2, pp. 870–878, 2014.

[12] J. Gao, E. C. Kara, S. Giri, and M. Bergés, "A feasibility study of automated plug-load identification from high-frequency measurements," in *Signal and Information Processing (GlobalSIP), 2015 IEEE Global Conference on.* IEEE, 2015, pp. 220–224.

[13] L. Du, D. He, R. G. Harley, and T. G. Habetler, "Electric load classification by binary voltage–current trajectory mapping," *IEEE Transactions on Smart Grid*, vol. 7, no. 1, pp. 358–365, 2016.

[14] N. Sadeghianpourhamami, J. Ruyssinck, D. Deschrijver, T. Dhaene, and C. Develder, "Comprehensive feature selection for appliance classification in nilm," *Energy and Buildings*, vol. 151, pp. 98–106, 2017.

[15] J. M. Carlo, M. S. Barbeitos, and H. R. Lasker, "Quantifying complex shapes: elliptical fourier analysis of octocoral sclerites," *The Biological Bulletin*, vol. 220, no. 3, pp. 224–237, 2011.

[16] Ø. D. Trier, A. K. Jain, and T. Taxt, "Feature extraction methods for character recognition-a survey," *Pattern recognition*, vol. 29, no. 4, pp. 641–662, 1996.

[17] M. M. Deshmukh and M. Dhopeshwarkar, "3d face recognition using contour, fourier descriptors."

[18] W. Lorensen and H. E. Cline, "Marching cubes: A high resolution 3d surface construction algorithm," *Computer Graphics (SIGGRAPH 87 Proceedings)*, pp. 163–170, 1987.

[19] R. Szeliski, *Computer vision: algorithms and applications.* Springer Science & Business Media, 2010.

[20] F. P. Kuhl and C. R. Giardina, "Elliptic fourier features of a closed contour," *Computer graphics and image processing*, vol. 18, no. 3, pp. 236–258, 1982.

[21] J. Caple, J. Byrd, and C. N. Stephan, "Elliptical fourier analysis: fundamentals, applications, and value for forensic anthropology," *International Journal of Legal Medicine*, pp. 1–16, 2017.

[22] J. Gao, S. Giri, E. C. Kara, and M. Bergés, "Plaid: a public dataset of high-resolution electrical appliance measurements for load identification research: demo abstract," in *proceedings of the 1st ACM Conference on Embedded Systems for Energy-Efficient Buildings.* ACM, 2014, pp. 198–199.