

Neural Network Ensembles to Real-time Identification of Plug-level Appliance Measurements

Karim Said Barsim, Lukas Mauch, and Bin Yang

Institute of Signal Processing and System Theory, University of Stuttgart

{karim.barsim, lukas.mauch, bin.yang}@iss.uni-stuttgart.de

Abstract— The problem of identifying end-use electrical appliances from their individual consumption profiles, known as the appliance identification problem, is a primary stage in both Non-Intrusive Load Monitoring (NILM) and automated plug-wise metering. Therefore, appliance identification has received dedicated studies with various electric appliance signatures, classification models, and evaluation datasets. In this paper, we propose a neural network ensembles approach to address this problem using high resolution measurements. The models are trained on the raw current and voltage waveforms, and thus, eliminating the need for well engineered appliance signatures. We evaluate the proposed model on a publicly available appliance dataset from 55 residential buildings, 11 appliance categories, and over 1000 measurements. We further study the stability of the trained models with respect to training dataset, sampling frequency, and variations in the steady-state operation of appliances.

Index Terms— Plug-level monitoring, Distributed sensing, Neural network ensembles, PLAID dataset

I. INTRODUCTION

The energy end-use monitoring problem has received a widespread attention over the last two decades, either in a non-intrusive fashion using smart meters which aim at reducing the cost of installation and maintenance [1, 2] or distributed sensing and smart outlets that became available through recent technological developments [3, 4]. Both trends face what is referred to as the appliance identification problem. In non-intrusive monitoring, the disaggregation stage reconstructs end-appliance consumption profiles from the aggregate measurements in a building. Afterwards, the appliance identification stage assigns either a generic end-use category [5] or specific appliance instances [6] to each reconstructed load profile. In plug-level metering, a smart outlet monitors the real-time consumption of its plugged-in appliance but this appliance may change over time based on user requirements. Similarly, identifying the connected appliance from the given consumption profile without user intervention is the challenge of appliance identification.

Previous works have addressed the problem of appliance identification using various machine learning tools, electric signatures, or evaluation platforms. In [3] for instance, support vector machines (SVM) and decision trees (DTree) trained on three-month, second-based power profiles of over 20 appliances achieved comparable performance on future runs of the same appliances. The performance degraded notably as previously unseen appliances were added to the test. In [7], current harmonics in addition to a few more features were extracted from high-resolution current waveforms (during transient and steady-state operations) to compare various

classification algorithms. High accuracies were reported using Bayesian networks evaluated on 16 categories of appliances and over 3000 acquired measurements. A more recent study, and most related to this work, compared and evaluated nine classifiers (including kNNs, logistic regression, SVMs, and DTrees) on different raw and engineered features [4] from a publicly available power dataset [8]. The authors proposed a novel appliance signature based on quantization of the VI-trajectories into a binary VI-image. The study further examined the effect of the sampling frequency on the identification problem showing that high classification rates can be obtained using random forest trees trained on the raw current and voltage measurement of at least 4kHz of sampling frequency (up to 82% of accuracy using VI-images and 86% using a set of combined features).

In most of these works neural networks on raw measurements have received little or no attention as a candidate model for appliance identification. Neural networks have recently received a growing interest in various machine learning applications especially after the evaluation of a deep convolutional architecture on a visual recognition dataset known as ImageNet. Recently, different deep learning architectures were found to achieve comparable results in the problem of energy disaggregation [9, 10].

In this paper, we first complement the work in [4] with an ensemble of neural networks in addressing the appliance identification problem from the raw, high-resolution current and voltage waveforms. We further study the stability and robustness of the proposed model with respect to the size of the training dataset, the sampling frequency of the measurements, and signal variations during the steady-state operation of appliances. We evaluate our models on the publicly available plug load appliance identification dataset (PLAID) [8] which contains, at the time of this work, measurements from 11 appliance categories, 235 appliance instances distributed over 55 households, and a total of 1074 current and voltage measurements at a sampling frequency $f_s = 30\text{kHz}$ and a grid frequency $f_g = 60\text{Hz}$.

II. APPLIANCES' SIGNATURES

In this work, we utilize the raw current $i(n)$ and voltage $v(n)$ waveforms during the steady-state operation as an appliance signature. Unlike previous works on VI-trajectories that extracted shape-based features [11, 12], we exploit one of the advantages of neural nets by utilizing the raw signals directly.

Given current $i(n)$ and voltage $v(n)$ signals of an appliance at a sampling frequency f_s and a grid frequency f_g , we extract

one complete period of each waveform

$$\varphi_\tau = [\varphi(\tau), \varphi(\tau+1), \dots, \varphi(\tau+d-1)]^T \in \mathbb{R}^d, \varphi \in \{i, v\}$$

where $d = f_s/f_g$ is the number of samples per period and τ is a point in time during the steady-state operation of the appliance. In order to support category-based appliance identification, we alleviate all amplitude information through segment-based normalization where each extracted signal segment φ_τ is normalized to the range $[-1,1]$. The result is the signal segments

$$\hat{\varphi}_\tau = \frac{2\varphi_\tau - (\max \varphi_\tau + \min \varphi_\tau) \mathbf{J}_{d,1}}{(\max \varphi_\tau - \min \varphi_\tau)} \in [-1,1]^d \quad (1)$$

where $\mathbf{J}_{m,n}$ is an m -by- n matrix of ones. The input vector to the neural network becomes

$$\mathbf{x} = \left[\hat{\mathbf{i}}_\tau^T, \hat{\mathbf{v}}_\tau^T \right]^T \in [-1,1]^{2d} \quad (2)$$

which corresponds to an input layer of size $2d$ for each neural network. We further expand the size of the training set algorithmically through extracting several signal segments from equally spaced initial phases of each measurement

$$\mathcal{X} = \left\{ \mathbf{x}_\tau \mid \tau = \tau_o + i\varepsilon, 0 \leq i < \frac{d}{\varepsilon} \right\} \quad (3)$$

where ε is the sliding step of the extraction window and this is applied to each measurement. Algorithmic expansion of the training data has several advantages. It provides a larger set of training data, eliminates the need for phase alignment of the signals, and more importantly drives a fully connected neural network to become invariant to the initial phase of the extracted signals. As a result, the model becomes more robust to signal variations during steady-state operation of appliances.

III. PREDICTION MODEL

Given a set $\Omega = \{\omega_m\}_{m=1}^M$ of M class labels (i.e. appliance categories), the straightforward approach for appliance classification is a multi-class neural network. However, in our initial tests we observed that an ensemble of binary networks performs consistently better than the multi-class model with a similar total size (i.e. neurons), training algorithm, regularization method, and overall training time. It is known that neural networks are one of the unstable learning algorithms with respect to training data [13] (i.e. they are sensitive to changes in the training set). For unstable models, an ensemble approach known as Bootstrap aggregation (aka Bagging) [14, 15] is expected to provide better results than a single base model¹. We, therefore, adopt the latter approach (i.e. an ensemble of binary networks) and discuss it further in the following.

For each class combination $(\omega_i, \omega_j)_{i < j}$, a binary classification neural network $\hat{\theta}_{\omega_i, \omega_j}$ is trained on an appropriate subset of the training data

$$\mathcal{D}_{\omega_i, \omega_j} = \{ \mathbf{x} \in \mathcal{X} \mid \omega(\mathbf{x}) = \omega_i \text{ or } \omega(\mathbf{x}) = \omega_j \} \quad \forall (i < j)$$

where $\omega(\mathbf{x})$ is the true class of the sample \mathbf{x} . Therefore, we train a total of $\binom{M}{2}$ base models each can be described as

¹In Bootstrap aggregation, an ensemble of weak, multi-class learners, each trained on a randomly sampled subset of the whole training dataset, is used rather than binary-classifiers generalized to a multi-classification problem.

$$\hat{\theta}_{\omega_i, \omega_j} : \mathcal{X} \mapsto [0,1]^2 \quad \forall (i < j) \quad (4)$$

$$\hat{\theta}_{\omega_i, \omega_j}(\mathbf{x}) = [\hat{p}_{\omega_i, \omega_j}(\mathbf{x}), \hat{p}_{\omega_j, \omega_i}(\mathbf{x})] \quad (5)$$

where $\hat{p}_{\omega_i, \omega_j}$ is an estimated normalized score for the class ω_i over ω_j for the given sample \mathbf{x} . Assuming a normalized output of the network

$$\hat{p}_{\omega_i, \omega_j}(\mathbf{x}) = 1 - \hat{p}_{\omega_j, \omega_i}(\mathbf{x}) \quad \forall \mathbf{x}, \quad (6)$$

the target vector for the training of $\hat{\theta}_{\omega_i, \omega_j}(\mathbf{x})$ becomes $[1,0]$ if $\omega(\mathbf{x}) = \omega_i$ or $[0,1]$ otherwise. The final ensemble prediction function is defined as

$$\hat{\omega}(\mathbf{x}) = \arg \max_{\omega_i \in \Omega} \sum_{j \neq i} \hat{p}_{\omega_i, \omega_j}(\mathbf{x}) \quad (7)$$

for a confidence-weighted voting or

$$\hat{\omega}(\mathbf{x}) = \arg \max_{\omega_i \in \Omega} \sum_{j \neq i} \mathbf{1}(\hat{p}_{\omega_i, \omega_j}(\mathbf{x}) > \hat{p}_{\omega_j, \omega_i}(\mathbf{x})) \quad (8)$$

for unweighted majority voting where $\mathbf{1}$ is the indicator function.

Finally, we utilize a validation-based early stopping as a regularization technique to avoid over fitting [16]. As mentioned earlier, PLAID measurements contains, for each appliance category, several appliance instances from various households. In order to improve generalization ability to new buildings, a 30% validation set is selected in a building-based fashion.

IV. EVALUATION

We evaluate our models on the publicly available PLAID dataset [8] which contains $M = 11$ appliance categories. Therefore, we have $\binom{M}{2} = 55$ class combinations and for each we train a two-layer, fully connected, feed-forward neural network with $2d$ input neurons, 30 hidden neurons, and two output neurons. Since some measurements contain turn-on transients, which are to be avoided in this work, the training set is extracted from the last two periods of each measurement. With a sliding step of $\varepsilon = 10$ samples, the training dataset is expanded by a factor of $d/\varepsilon = 50$. The hyperbolic tangent (*tanh*) activation is utilized in the hidden layer while the normalized exponential (*softmax*) is selected for the output layer. We utilized `matlab`'s implementation of the conjugate gradient descent with random restarts [17] as a training function for all models.

In order to compare this approach with the previously studied models on this dataset [4], we adopt the same *leave-house-out* cross validation technique. In other words, samples from one house is saved for test while the remaining set of buildings are available for training, resulting in a total of 55 test cases.

For each class ω_m and the total confusion matrix $\Lambda = [\Lambda_{ij}]_{i,j=1}^M$, we introduce the following performance indicators

$$\text{TP}_m = \left| \left\{ \mathbf{x} \in \tilde{\mathcal{X}} \mid \hat{\omega}(\mathbf{x}) = \omega_m \ \& \ \omega(\mathbf{x}) = \omega_m \right\} \right| \quad (9)$$

$$\text{TN}_m = \left| \left\{ \mathbf{x} \in \tilde{\mathcal{X}} \mid \hat{\omega}(\mathbf{x}) \neq \omega_m \ \& \ \omega(\mathbf{x}) \neq \omega_m \right\} \right| \quad (10)$$

$$\text{FP}_m = \left| \left\{ \mathbf{x} \in \tilde{\mathcal{X}} \mid \hat{\omega}(\mathbf{x}) = \omega_m \ \& \ \omega(\mathbf{x}) \neq \omega_m \right\} \right| \quad (11)$$

$$\text{FN}_m = \left| \left\{ \mathbf{x} \in \tilde{\mathcal{X}} \mid \hat{\omega}(\mathbf{x}) \neq \omega_m \ \& \ \omega(\mathbf{x}) = \omega_m \right\} \right| \quad (12)$$

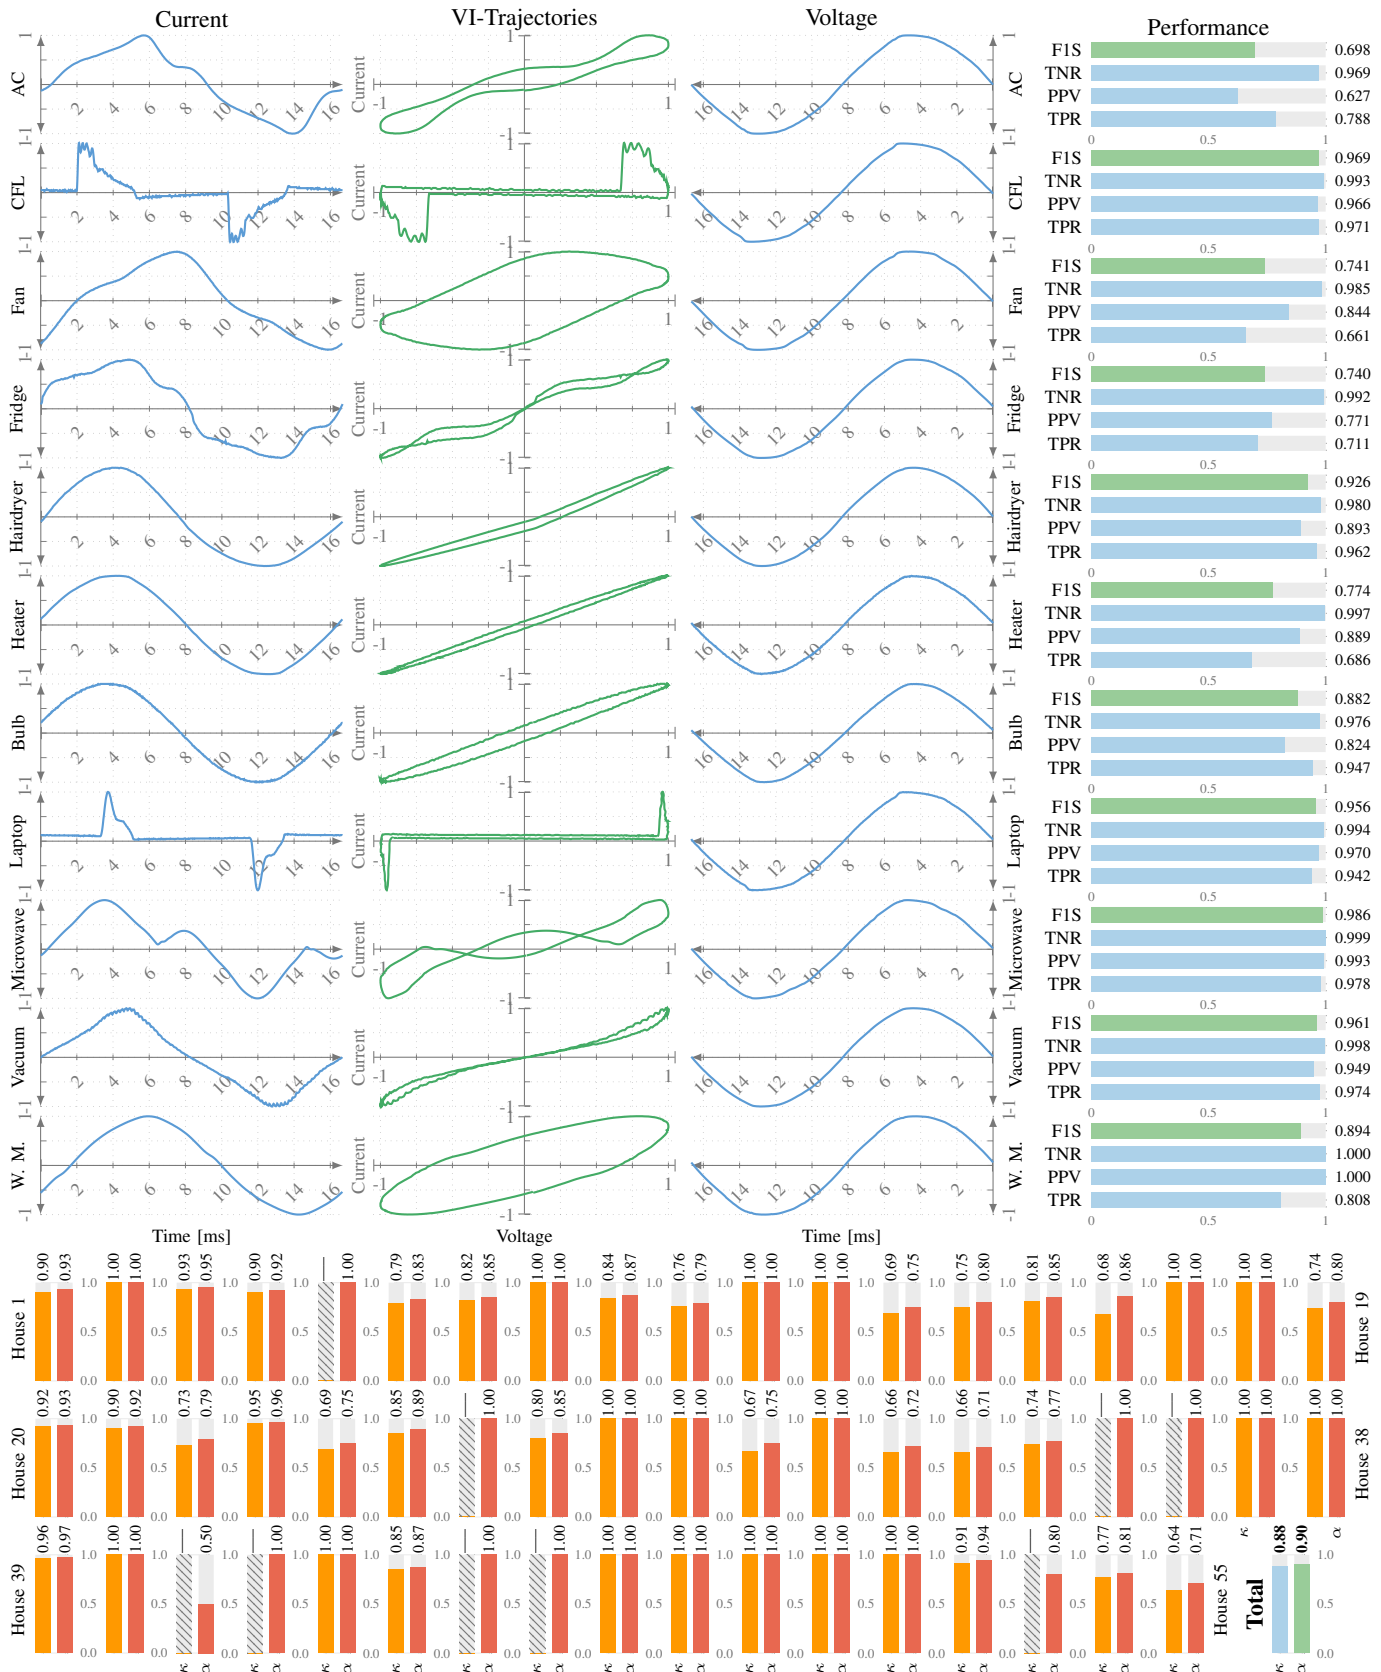


Fig. 1: Current, voltage, and VI-trajectories of selected samples from PLAID dataset [8] are shown to the left. The appliance categories and IDs, from top to bottom, are (1) air conditioner, 1010, (2) CFL, 20, (3) fan, 766, (4) fridge, 6, (5) hairdryer, 444, (6) heater, 716, (7) bulb, 57, (8) laptop 28, (9) microwave, 10, (10) vacuum cleaner, 730, and (11) washing machine, 488. The per-category evaluation metrics are shown to the right while the per-house metrics are in the bottom. The best evaluation results are $\kappa = 0.882$ and $\alpha = 0.897$.

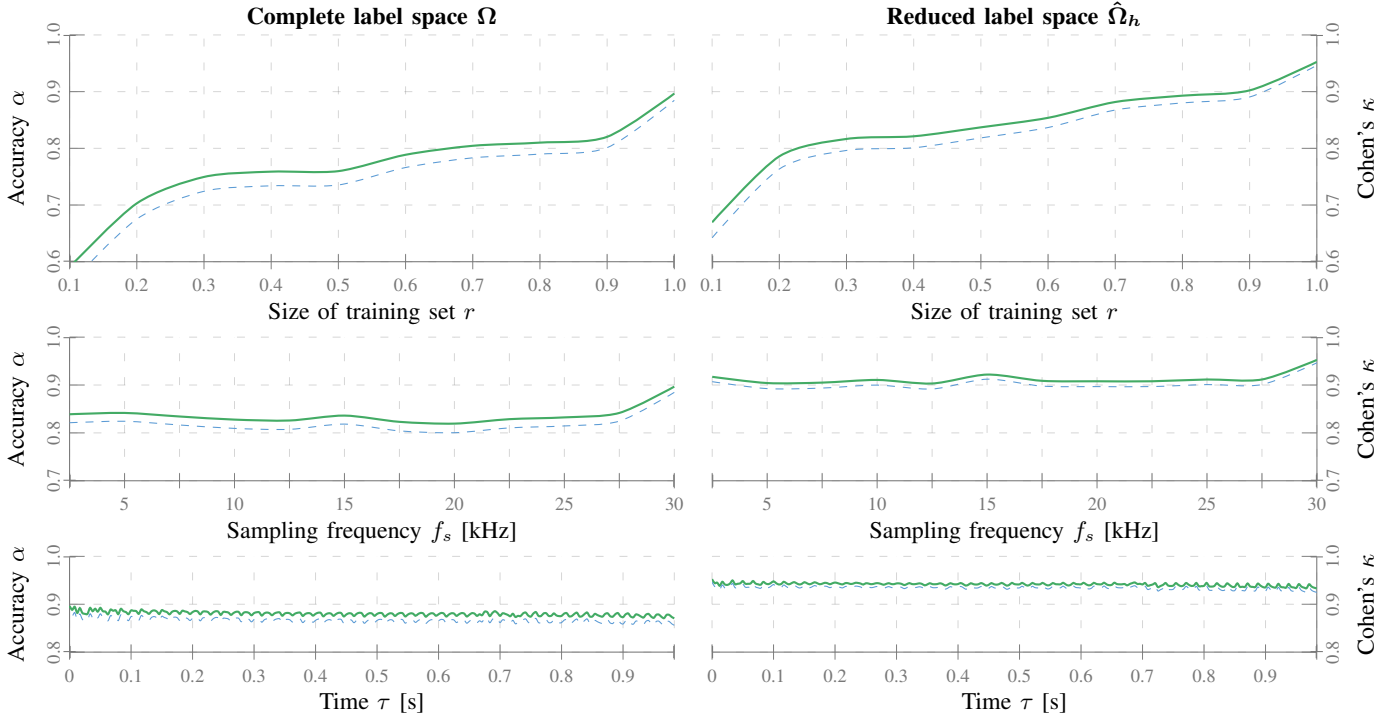


Fig. 2: Aggregate evaluation results as a function of (a) training set size, (b) sampling frequency, and (c) time shift. To right are the same experiments with reduced label spaces based on the *prior knowledge* of the label space of each target household.

$$\text{recall}_m = \text{TPR}_m = \frac{\text{TP}_m}{\text{TP}_m + \text{FN}_m} \quad (13)$$

$$\text{precision}_m = \text{PPV}_m = \frac{\text{TP}_m}{\text{TP}_m + \text{FP}_m} \quad (14)$$

$$\text{specificity}_m = \text{TNR}_m = \frac{\text{TN}_m}{\text{TN}_m + \text{FP}_m} \quad (15)$$

$$F_1^m\text{-score} = \text{FIS}_m = \frac{2\text{TP}_m}{2\text{TP}_m + \text{FP}_m + \text{FN}_m} \quad (16)$$

for per-class evaluation whereas the unweighted accuracy α and Cohen's κ

$$\alpha = \frac{\text{tr}(\mathbf{\Lambda})}{\sum_{i,j=1}^M \Lambda_{i,j}} = \frac{\text{tr}(\mathbf{\Lambda})}{\mathbf{J}_{1,M} \mathbf{\Lambda} \mathbf{J}_{M,1}}, \quad (17)$$

$$\kappa = \frac{\text{tr}(\mathbf{\Lambda}) \times \mathbf{J}_{1,M} \mathbf{\Lambda} \mathbf{J}_{M,1} - \text{tr}(\mathbf{\Lambda} \mathbf{J}_M \mathbf{\Lambda})}{(\mathbf{J}_{1,M} \mathbf{\Lambda} \mathbf{J}_{M,1})^2 - \text{tr}(\mathbf{\Lambda} \mathbf{J}_M \mathbf{\Lambda})} \quad (18)$$

are utilized for the aggregate, multi-class evaluation.

Figure 1 shows the per-class, and per-household performance results. The best result obtained for the voting in Equation 7 is $\alpha = 89.7\%$ while the unweighted majority vote Equation 8 achieved 88.2% of accuracy. It is observed from the figures that microwaves, compact fluorescent lamps, vacuum cleaners, and laptops are the most identifiable loads whereas temperature control devices (such as air conditioners, heaters, fridges, and even fans) are the least.

Figure 2 shows further tests of stability of the adopted model. In the first study, we reduce the size of the training dataset and estimate the total performance using the same cross validation technique. In each test case, one building is saved for evaluation, $54r$ buildings for training where $0 < r \leq 1$ (which is further partitioned into training and validation households), and the remaining $54(1 - r)$ are untouched. As expected, the performance of the model degrades notably as the size of the training data decreases.

In the second study, we use the complete training dataset (i.e. $r = 1$) but reduce the sampling frequency of the raw mea-

surements. A suitable FIR filter is utilized for each fractional decimation. This study reveals the robustness of our method with respect to the sampling frequency. As observed in the figure, the accuracy slightly drops but is always above 80% even at 2.5kHz.

In the third study, we test the sensitivity of the model to the phase shift of the extracted segments. In other words, a single period from each measurement with phase τ is extracted for testing. The common shortest period of all PLAID measurements is one second. Similarly, the figure reveals the robustness of the proposed model to signal variations in steady-state operation, at least for the short common period available in PLAID.

Finally, we repeat all previous experiments while exploiting external knowledge about the list of appliances in each house in reducing the label space for each test case. In other words, if house $h = 1$ does not have an air conditioner, the label space for this building becomes $\hat{\Omega}_1 = \Omega - \{\text{'AC'}\}$ and we rather train 45 binary networks. This of course improves the performance where the best-case accuracy reaches $\alpha = 94\%$, but the question is of course about the availability of this knowledge in practice.

V. CONCLUSION AND FUTURE WORK

In this work, we introduced an approach of ensemble of neural networks in addressing the appliance identification problem using raw, high-resolution current and voltage measurements. We evaluated our model on the PLAID dataset with the best performance test achieving 89% of accuracy.

It was observed that the performance notably degrades as the size of the training data is reduced. Therefore, one of the main future work packages is semi-supervised learning methods.

REFERENCES

- [1] G. W. Hart, "Nonintrusive appliance load monitoring". In *proceedings of the IEEE* 80 (12):1870–1891, Dec. 1992.
- [2] J. Froehlich, E. Larson, S. Gupta, G. Cohn, M. S. Reynolds, and S. N. Patel, "Disaggregated End-Use Energy Sensing for the Smart Grid". In *proceedings of Pervasive Computing, IEEE* 10 (1):28–39, Jan. 2011.
- [3] S. Barker, M. Musthag, D. Irwin, and P. Shenoy, "Non-Intrusive Load Identification for Smart Outlets". In *proceedings of the IEEE International Conference on Smart Grid Communications (SmartGridComm)*:548–553, Nov. 2014.
- [4] J. Gao, E. C. Kara, S. Giri, and M. Bergés, "A feasibility study of automated plug-load identification from high-frequency measurements". In *proceedings of the 3rd IEEE Global Conference on Signal and Information Processing (GlobalSIP)*:220–224, Dec. 2015.
- [5] O. Parson, S. Ghosh, M. Weal, and A. Rogers, "An Un-supervised Training Method for Non-intrusive Appliance Load Monitoring". *Artificial Intelligence* 217:1–19, 2014.
- [6] J. Z. Kolter and T. Jaakkola, "Approximate Inference in Additive Factorial HMMs with Application to Energy Disaggregation". In *proceedings of the International Conference on Artificial Intelligence and Statistics* 22:1472–1482, 2012.
- [7] A. Reinhardt, D. Burkhardt, M. Zaheer, and R. Steinmetz, "Electric Appliance Classification Based on Distributed High Resolution Current Sensing". In *proceedings of the 37th IEEE Conference on Local Computer Networks (LCN)*:999–1005, Oct. 2012.
- [8] 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*:198–199, 2014.
- [9] J. Kelly and W. J. Knottenbelt, "Neural NILM: Deep Neural Networks Applied to Energy Disaggregation". *CoRR* abs/1507.06594, Aug. 2015. [Online]. Available: <http://arxiv.org/abs/1507.06594>
- [10] L. Mauch and B. Yang, "A new approach for supervised power disaggregation by using a deep recurrent LSTM network". In *proceedings of the 3rd IEEE Global Conference on Signal and Information Processing (GlobalSIP)*:63–67, Dec. 2015.
- [11] H. Y. Lam, G. Fung, and W. K. Lee, "A Novel Method to Construct Taxonomy Electrical Appliances Based on Load Signatures". *Consumer Electronics, IEEE Transactions on* 53 (2):653–660, May 2007.
- [12] T. Hassan, F. Javed, and N. Arshad, "An Empirical Investigation of V-I Trajectory Based Load Signatures for Non-Intrusive Load Monitoring". *IEEE Transactions on Smart Grid* 5 (2):870–878, Mar. 2014.
- [13] C. C. Aggarwal, *Data Classification Algorithms and Applications*. Chapman & Hall/CRC, 2015.
- [14] L. Breiman, "Bagging Predictors". *Machine Learning* 24 (2):123–140, 1996.
- [15] J. H. Friedman and P. Hall, "On bagging and nonlinear estimation". *Journal of Statistical Planning and Inference* 137 (3):669 – 683, 2007.
- [16] L. Prechelt, "Early Stopping - but when?". *Neural Networks: Tricks of the Trade*:53–67, 2012.
- [17] M. J. D. Powell, "Restart procedures for the conjugate gradient method". *Mathematical Programming* 12 (1):241–254, 1977.