A Scalable Non-intrusive Load Monitoring System for Fridge-Freezer Energy Efficiency Estimation

Oliver Parson, Mark Weal, Alex Rogers
Electronics and Computer Science
University of Southampton, Hampshire, SO17 1BJ, UK
{osp,mjw,acr}@ecs.soton.ac.uk

***Abstract*—**
In this paper we propose an approach by which the energy efficiency of individual appliances can be estimated from an aggregate load. To date, energy disaggregation research has presented results for small data sets of 7 households or less, and as a result the generality of results are often unknown. In contrast, we have deployed household electricity sensors to 117 households and evaluated the accuracy by which our approach can identify the energy efficiency of refrigerators and freezers from an aggregate load. Crucially, our approach does not require training data to be collected by sub-metering individual appliances, nor does it assume any knowledge of the appliances present in the household. Instead, our approach uses prior models of general appliance types that are used to first identify which households contain either a combined fridge-freezer or separate refrigerator and freezer, and subsequently to estimate the energy efficiency of such appliances. Finally, we calculate the time until the energy savings of replacing such appliances have offset the cost of the replacement appliance, which we show can be as low as 2.5 years.

1. Introduction

Non-intrusive appliance load monitoring (NIALM), or energy disaggregation, aims to break down a household's aggregate electricity consumption into individual appliances [1]. The motivation for such technology is to be able to inform a household's occupants of how much energy each appliance consumes and furthermore offer actionable energy saving advice to empower them to reduce their consumption [2]. To realise this application through a practical and widely applicable software system, it is essential to take advantage of existing infrastructure rather than designing new hardware. Smart meters are currently being deployed on national scales [3] and will report household aggregate power at 10 second intervals, and are therefore an ideal data collection platform for NIALM.

Recent contributions to the field of NIALM have applied principled machine learning techniques to the problem of energy disaggregation. However, all approaches applicable to 10 second smart meter data have only been applied to data sets of 7 households or less, and furthermore such test households often contain an unrealistically small number of appliances. As a result, it is not known whether such approaches are either over-fitted to the test households or are sufficiently general to be applied at larger scales. For example, Kolter & Johnson proposed an approach based on the factorial hidden Markov model (FHMM), however it was only tested in 5 households [4]. Kim et al. proposed an unsupervised approach also based on FHMMs, which was only tested in 7 households, each of which contained a maximum of 10 appliances [5]. Parson et al. proposed an approach which incorporates prior knowledge into HMM-based disaggregation, but was only tested on 4 households [6]. Most recently, Johnson & Willsky proposed an approach based on the factorial hidden semi-Markov model (FHSMM), yet it was only tested in 4 households, each of which contained a maximum of 5 appliances [7]. Clearly, there exists a challenge to evaluate NIALM approaches in larger numbers of households before conclusions can be drawn regarding their generality and scalability.

In this paper, we address the challenge of evaluating the scalability and generality of energy disaggregation, and to do so, we present a deployment of an extension of our previously proposed approach [6] to 117 real UK households. In each household, we install an electricity sensor which reports the household aggregate power demand at 10 second intervals, matching the data that will be reported by UK smart meters [3]. We then show how our approach is able to estimate the energy efficiency of the refrigerators and freezers in these households without requiring sub-metered appliance data to be collected or information to be collected regarding the appliances present in the household. Crucially, this requires our approach to first determine which households contain either a combined fridge-freezer or a separate refrigerator and freezer. Next, our approach identifies durations when only the refrigerator and freezer are operating (typically during overnight periods), from which it learns a model of the appliances' behaviour. From this model, the energy efficiency of the refrigerator and freezer are calculated and compared to an energy efficient replacement. We show that our approach is able to correctly detect 86% of households which contain a combined fridge-freezer, while producing less than 20% false positives. We then demonstrate that our approach is able to estimate the energy efficiency of such fridge-freezers with an average error of 11%. Next, we show that this information can be used to calculate the time until the energy savings due to replacing the appliance have offset the cost of the new appliance, which can be as low as 2.5 years. We also provide similar results for households with separate refrigerators and freezers.

2. Problem Description

The aim of appliance energy efficiency estimation is as follows. Given a discrete sequence of observed aggregate power readings, $\bar{\mathbf{y}} = \bar{y}_1, \ldots, \bar{y}_T$, determine an appliance's annual energy consumption, $E^{(n)}$, where n is one of N appliances. In the field of energy disaggregation, this often involves modelling the power demand, $\mathbf{y}^{(n)} = y_1^{(n)}, \ldots, y_T^{(n)}$, and operating state, $\mathbf{x}^{(n)} = x_1^{(n)}, \ldots, x_T^{(n)}$, of appliance n in time slice t.

3. Deployment of Household Electricity Sensors

We deployed household electricity sensors to 117 households in the village of Colden Common, Hampshire, UK. We attached an AlertMe (www.alertme.com) current transformer clamp to the household electricity input, shown by Figure 1 (a). The current transformer measures the electromagnetic field surrounding the wire, from which the household power demand can be calculated. As such, the sensor captures the same data that will be reported by UK smart meters. The power data is then transmitted to an in-home hub at 10 second intervals via a ZigBee network, before being uploaded to a cloud storage via an Ethernet connection to each home's router.

4. Appliance Energy Estimation

In this section, we describe our approach which is able to estimate the energy efficiency of both combined fridge-freezers and separate refrigerator and freezers given only the household aggregate power demand. First, we present the probabilistic graphical model which we use to represent appliances. Second, we describe the method by which the model can be used to identify individual appliance's signatures in the household aggregate load. Third, we show how the extracted signatures can be used to update the appliance model parameters to represent the behaviour of the specific appliance instances in a single household. Last, we show how such models can be used to estimate the efficiency of that household's appliances.

A. Appliance Modelling via Bayesian Hidden Markov Models

We adopt a state-of-the-art Bayesian approach to appliance modelling using hidden Markov models (HMMs) [6, 7], as shown by Figure 1 (b). In such a model, the discrete variables of the model, $\mathbf{x}$, represent the states of an appliance, while the continuous variables, $\mathbf{y}$, represent the power demand of that appliance. The start of the chain of appliance states is distributed according to a categorical distribution with parameters $\boldsymbol{\pi}$, and a prior Dirichlet distribution with hyperparameters $\hat{\boldsymbol{\alpha}}$. The appliance's state evolves via Markovian dynamics according to a transition matrix, $\mathbf{A}$, for which the prior distribution is a vector of Dirichlet distributions with hyperparameters $\hat{\mathbf{C}}$. Last, we model the power demand of each appliance state using a Gaussian distribution with mean $\boldsymbol{\mu}$ and precision $\boldsymbol{\tau}$, for which the prior distribution is a Gaussian-Gamma distribution, with mean $\hat{\boldsymbol{\lambda}}$ and precision $\hat{\mathbf{r}}$, and shape $\hat{\boldsymbol{\beta}}$ and scale $\hat{\mathbf{w}}$, respectively.

The hyperparameters of the model, $\hat{\boldsymbol{\theta}} = \{\hat{\boldsymbol{\alpha}}, \hat{\mathbf{C}}, \hat{\boldsymbol{\lambda}}, \hat{\mathbf{r}}, \hat{\boldsymbol{\beta}}, \hat{\mathbf{w}}\}$, represent the behaviour of a broad appliance type (e.g. all refrigerators), and are learned from the Tracebase appliance data set [8]. However, the appliance parameters, $\boldsymbol{\theta} = \{\boldsymbol{\pi}, \mathbf{A}, \boldsymbol{\mu}, \boldsymbol{\tau}\}$, represent the behaviour of a specific appliance instance in a single household, and are therefore the variables which we aim to infer from household aggregate data.

We use this HMM to represent the behaviour of combined fridge-freezers within a constant aggregate load. However, in households with a separate refrigerator and freezer, each such appliance will produce a separate signature. As a result, the sum of the two signals will appear within the aggregate load. In such cases, we use a factorial hidden Markov model (FHMM) which consists of two chains of discrete variables corresponding to the separate refrigerator and freezer.

B. Identifying Appliance Signatures

We use the probabilistic appliance model described in the previous section to search for periods of the household aggregate power demand during which only either the combined fridge-freezer or separate refrigerator and freezer are changing state, as proposed in our previous work [6]. This is achieved by sliding a window through the aggregate data and calculating the likelihood that the window of data was generated by the modelled appliances. We first apply the basic HMM, and if at least one signature is found that household is classified as containing a combined fridge-freezer. If no signature is found, we apply the FHMM with two appliance chains, and if at least one signature is found that household is classified as containing a separate refrigerator and freezer. However, if no signature is found by either model, that household is classified as containing either zero or three or more refrigerators and freezers. Figure 1 (c) shows an example of this method for a household containing a combined fridge-freezer, from which it can be seen that the 02:00 - 05:00 window scores a high likelihood of being generated by the fridge-freezer model.

However, it is important to note that our approach aims to extract periods during which only the appliances of interest are changing state, and that other appliances might be drawing a constant power during this period. Therefore, in our approach, the base-load is first subtracted from the aggregate load:

$$\tilde{\mathbf{y}}_{i:j} = \bar{\mathbf{y}}_{i:j} - \min(\bar{\mathbf{y}}_{i:j}) \tag{1}$$

where $\bar{\mathbf{y}}_{i:j}$ is a window of aggregate data $\bar{\mathbf{y}}_i, \ldots, \bar{\mathbf{y}}_j$, and $\tilde{\mathbf{y}}_{i:j}$ is the same window after the base-load has been subtracted. This ensures that the distributions over each state's mean power demand correspond between signatures. We determine whether power data was generated only by the modelled appliances:

$$accept(\tilde{\mathbf{y}}_{i:j}) = \begin{cases} true & \text{if } p(\tilde{\mathbf{y}}_{i:j}|\hat{\boldsymbol{\theta}}) > D \\ false & \text{otherwise.} \end{cases} \tag{2}$$

where $\tilde{\mathbf{y}}_{i:j}$ is a window of aggregate data after the base-load has been subtracted, D is an appliance specific likelihood threshold, and $p(\tilde{\mathbf{y}}_{i:j}|\hat{\boldsymbol{\theta}})$ is the likelihood of that window of data given the general appliance model as calculated by:

$$p(\tilde{\mathbf{y}}_{i:j}|\hat{\boldsymbol{\theta}}) = \iint p(\tilde{\mathbf{y}}_{i:j}, \mathrm{x}|\boldsymbol{\theta}) p(\boldsymbol{\theta}|\hat{\boldsymbol{\theta}}) \, \mathrm{dx} \, \mathrm{d}\boldsymbol{\theta} \tag{3}$$

which intuitively represents the likelihood of the data, $\tilde{\mathbf{y}}_{i:j}$, given a general appliance model, $\hat{\boldsymbol{\theta}}$, with the unknown states

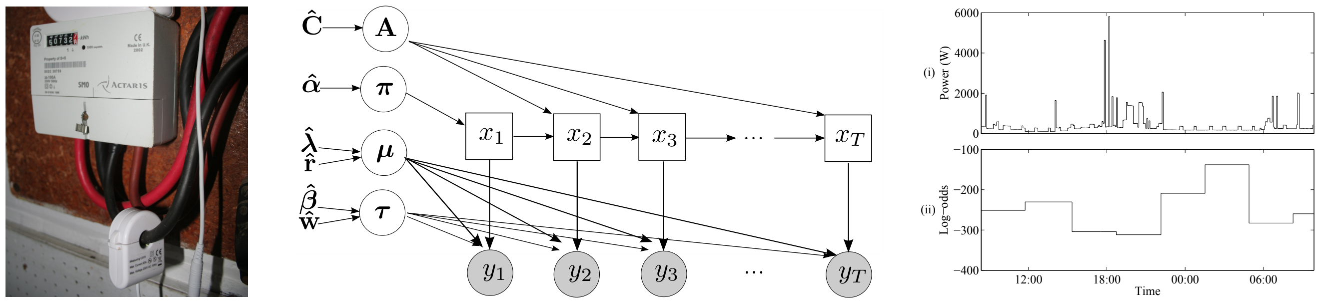

(a) Household aggregate electricity sensor

(b) Detection accuracy of households with separate refrigerators and freezers.

(c) (i) Household power demand over 24 hours. (ii) Log-odds of power with fridge-freezer model.

Fig. 1: Our approach to model appliances from single point of measurement.

$\mathbf{x}$ and instance parameters $\boldsymbol{\theta}$ of the appliances integrated out. The threshold D is set such that the model will accept windows of data which can be sufficiently explained by either the HMM or FHMM model, and reject any windows of data generated by other appliance types or combinations of appliances. This process identifies windows of aggregate data in which only one or two appliances matching the general model change state.

C. Inferring Appliance Model Parameters

Our approach uses the signatures extracted from the household aggregate load to update the general appliance parameters to represent the behaviour of the specific appliances in a single household. For households with a combined fridge-freezer, we infer a posterior distribution over the variables $\boldsymbol{\theta}$ via variational message passing [9] implemented using Infer.NET [10]. However, Infer.NET does not support FHMMs, and therefore we use Gibbs sampling for inference using the `pyhsmm` Python library [7] for households with a separate fridge and freezer.

D. Estimating Appliance Energy Efficiency

Having inferred posterior distributions over the parameters of the appliances in a single household, our approach subsequently estimates the annual appliance energy consumption:

$$E^{(n)} = \frac{24 \times 365}{1000} \sum_{k=1}^{K} \mu_k^{(n)} \frac{A_{k,k}^{(n)}}{\text{trace}(\mathbf{A}^{(n)})} \tag{4}$$

where trace($\mathbf{A}$) is a function which returns the sum of diagonal elements of matrix $\mathbf{A}$.

5. Evaluation of Energy Efficiency Estimation

In this section, we first evaluate the accuracy by which our approach is able to detect whether each of the 117 households contains a combined fridge-freezer or an individual refrigerator and freezer. We then evaluate the accuracy by which our approach estimated the energy efficiency of such appliances.

We evaluate the classification accuracy of households using a ROC curve, which represents the trade-off between the fraction of households containing a certain appliance which were correctly detected, or true positive rate (TPR), and the fraction of households which did not contain a certain appliance but were incorrectly detected, or false positive rate (FPR), for various likelihood thresholds, D. Since existing approaches are only applicable to households for which models are known for all appliances, instead we compare the performance of our approach under two different general appliance models: a refrigerator appliance model as learned from the Tracebase data set [8] and a variant of this model in which the power demand of the *on* state matches that of a combined fridge-freezer. We also indicate the accuracy of a random classifier.

Figure 2 (a) shows a ROC curve which represents the detection accuracy of our approach for households with a combined fridge-freezer. It can be seen that the fridge-freezer general appliance model performs preferably to the refrigerator model for almost all trade-offs between TPR and FPR. This matches the intuition that the combined fridge freezer model should outperform a model representing only refrigerators for households which contain a combined fridge freezer. A particularly favourable trade-off is highlighted, at which the combined fridge-freezer model is able to correctly detect 86% of households while also producing less than 20% false positives. Figure 2 (b) shows a second ROC curve for households with a separate refrigerator and freezer. Unlike the results for the combined fridge-freezer model, it can be seen that neither approach appears to be preferable for a range of likelihood thresholds. This is likely due to the mixture of refrigerators, freezers and combined fridge-freezers in such households, and as a result, neither model produces superior performance.

Having evaluated the accuracy by which our approach is able to classify whether households contain either a combined fridge-freezer or a separate refrigerator and freezer, we now evaluate the accuracy by which our approach is able to estimate the energy efficiency of such appliances. We present results for the households which were correctly classified as containing either a combined fridge-freezer or a separate refrigerator and freezer. This represents a subset of the 117 households, as some households contained zero or more than two fridges and freezers while others were incorrectly classified as discussed above. We calculated the ground truth annual energy consumption by manually labelling individual appliance signatures within the aggregate load and scaling the energy consumption to one year.

Figure 2 (c) shows a histogram of the difference between the estimated and actual annual energy consumption of the

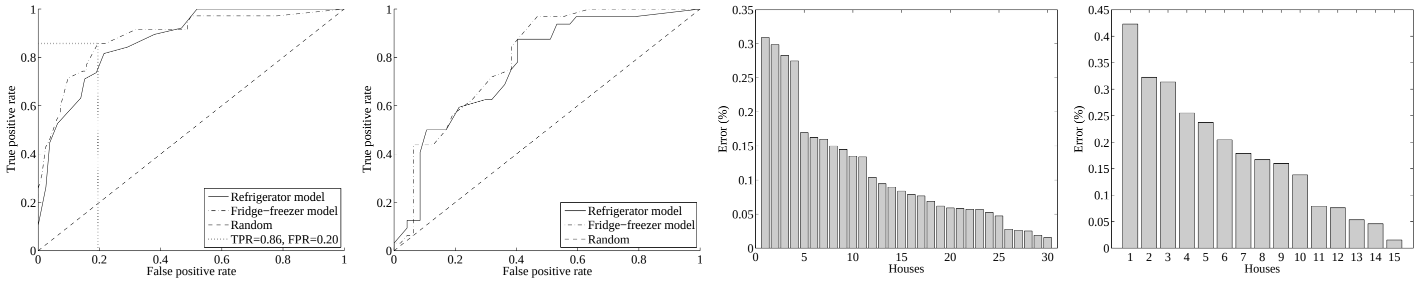

(a) Detection accuracy of households with combined fridge-freezers.

(b) Detection accuracy of households with separate refrigerators and freezers.

(c) Percentage error in estimate of annual fridge-freezer energy consumption.

(d) Percentage error in estimate of fridge and freezer energy consumption.

Fig. 2: Results of appliance detection and energy consumption estimation.

fridge-freezers. For each household, the error is normalised by the actual energy consumption of the fridge-freezer. It can be seen that our approach is able to estimate the energy consumption of the fridge-freezers to within 20% of the actual energy consumption for 26 of the 30 households, while the fridge-freezer energy consumption for the remaining 4 households are estimated with less than 35% error (mean = 11%). Therefore, it would be appropriate to provide feedback regarding the benefits of replacing their fridge-freezer to the households for which the fridge-freezer energy consumption was estimated to be more than 35% greater than that of an energy efficient replacement, which corresponds to 21 out of the 30 households. Furthermore, we calculated time until the annual financial savings have offset the cost of the replacement appliance, which was as low as 2.5 years for the household with the least efficient appliances, assuming a cost of electricity of £0.15/kWh and a cost of replacement of £429.

Figure 2 (d) shows a similar histogram for households containing a separate refrigerator and freezer. It can be seen that the error is systematically greater than for households with combined fridge-freezers, with only 9 of the 15 households' total refrigerator and freezer energy consumption estimated with less than 20% error, while the remaining 6 households were estimated with a maximum of 45% error. This indicates that it is more difficult to estimate the energy consumption in households with a combined fridge-freezer than in households with a separate refrigerator and freezer. However, although the normalised error is greater, so is the difference in energy consumption between the estimated total of refrigerators and freezers and a replacement combined fridge-freezer. As a result, it would be appropriate to provide feedback regarding the benefits of replacing their separate refrigerators and freezers with a combined fridge-freezer to the households for which the combined fridge-freezer energy consumption was estimated to be more than 45% greater than that of an energy efficient replacement, which corresponds to 10 of the 15 households. We calculcated the time until the annual financial savings have offset the cost of the replacement appliance to be as low as 5 years for the household with the least efficient appliances.

6. Conclusion

In this paper, we have presented the first application of energy disaggregation research to a large number of households' smart meter data. Crucially, our approach does not require training data to be collected by sub-metering individual appliances, nor does it assume any knowledge of the appliances present in the household. Instead, the approach uses prior models of general appliance types used to first identify which appliances are present in a household, and subsequently estimate the energy efficiency of such appliances. We evaluated the accuracy with which our approach could classify the type of refrigerator and freezer in a household, and the accuracy by which the energy efficiency of such appliances could be estimated, and showed that the energy saved by replacing the appliance can offset the cost of replacement in as little as 2.5 years. Future work will consist of a user study which will assess whether such feedback motivates the household occupants to replace their appliances and whether such actions produced the forecast savings.

References

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

[2] S. Darby, "The Effectiveness of Feedback on Energy Consumption, A review for DEFRA of the literature on metering, billing and direct displays," University of Oxford, UK, Tech. Rep., 2006.

[3] Department of Energy & Climate Change, "Smart Metering Equipment Technical Specifications Version 2," UK, Tech. Rep., 2013.

[4] J. Z. Kolter and M. J. Johnson, "REDD: A Public Data Set for Energy Disaggregation Research," in *ACM Special Interest Group on Knowledge Discovery and Data Mining, workshop on Data Mining Applications in Sustainability*, San Diego, CA, USA, 2011.

[5] H. Kim, M. Marwah, M. Arlitt, G. Lyon, and J. Han, "Unsupervised Disaggregation of Low Frequency Power Measurements," in *Proceedings of the 11th SIAM International Conference on Data Mining*, Mesa, AZ, USA, 2011, pp. 747–758.

[6] O. Parson, S. Ghosh, M. Weal, and A. Rogers, "Non-intrusive load monitoring using prior models of general appliance types," in *Proceedings of the 26th AAAI Conference on Artificial Intelligence*, Toronto, ON, Canada, 2012, pp. 356–362.

[7] M. J. Johnson and A. S. Willsky, "Bayesian Nonparametric Hidden Semi-Markov Models," *Journal of Machine Learning Research*, vol. 14, pp. 673–701, 2013.

[8] A. Reinhardt, P. Bauman, D. Burgstahler, M. Hollick, H. Chonov, M. Werner, and R. Steinmetz, "On the Accuracy of Appliance Identification Based on Distributed Load Metering Data," in *Proceedings of the 2nd Conference on Sustainable Internet and ICT for Sustainability*, Pisa, Italy, 2012, pp. 1–9.

[9] J. Winn and C. M. Bishop, "Variational message passing," *Journal of Machine Learning Research*, vol. 6, no. 1, pp. 661–694, 2006.

[10] T. Minka, J. Winn, J. Guiver, and D. Knowles, "Infer.NET 2.5," Microsoft Research Cambridge, Tech. Rep., 2012.