Non-intrusive Load Monitoring using Prior Models of General Appliance Types

[Extended Abstract]

Oliver Parson, Siddhartha Ghosh, Mark Weal, Alex Rogers
Electronics and Computer Science
University of Southampton, Hampshire, SO17 1BJ, UK
{op106,sg2,mjw,acr}@ecs.soton.ac.uk

ABSTRACT

Non-intrusive appliance load monitoring is the process of disaggregating a household's total electricity consumption into its contributing appliances. In this paper we propose an approach by which individual appliances can be iteratively separated from an aggregate load. Unlike existing approaches, our approach does not require training data to be collected by sub-metering individual appliances, nor does it assume complete knowledge of the appliances present in the household. Instead, we propose an approach in which prior models of general appliance types are tuned to specific appliance instances using only signatures extracted from the aggregate load. The tuned appliance models are then used to estimate each appliance's load, which is subsequently subtracted from the aggregate load. This process is applied iteratively until all appliances for which prior behaviour models are known have been disaggregated. This paper summarises the contribution described in full in [9].

1. INTRODUCTION

Non-intrusive load monitoring (NIALM), or energy disaggregation, aims to break down a household's aggregate electricity consumption into individual appliances [4]. The motivations for such a process are twofold. First, informing a household's occupants of how much energy each appliance consumes empowers them to take steps towards reducing their energy consumption [1]. Second, if the NIALM system is able to determine the current time of use of each appliance, a recommender system would be able to inform a household's occupants of the potential savings through deferring appliance use to a time of day when electricity is either cheaper or has a lower carbon footprint.

To address these goals through a practical and widely applicable software system, it is essential to take advantage of existing infrastructure rather than designing new hardware. Smart meters and in-home displays are currently being deployed on national scales [2] and thus constitute ideal data collection and information display platforms respectively for NIALM solutions. However, smart meters are only likely to transmit low granularity data to in-home displays (every 5 seconds maximum in the UK [3]). Therefore, NIALM methods must be applicable to such low granularity aggregate data if they are to be applied within such a scenario.

Recent contributions to the field of NIALM have applied principled machine learning techniques to the problem of energy disaggregation. Such approaches fall into two categories. The first uses supervised methods which assume that sub-metered (ground truth) data is available for training before the disaggregation task is performed [6; 8]. This assumption dramatically increases the investment required to set up such a system, since in practice, installing submeters may be inconvenient or time consuming. The second uses unsupervised disaggregation methods [5; 7; 11] in which no prior knowledge of the appliances is assumed, but which often require appliances to be manually labelled after the disaggregation process or assume knowledge of the number of household appliances. Such approaches also typically ignore additional information that may be available regarding which appliances are likely to be present in a house or how such appliances are likely to behave.

While these assumptions are attractive from a machine learning perspective, they do not address the most likely real world applications of NIALM; where sub-metered data and complete knowledge of the appliance set is not available, but some prior information about some appliances might be known. This prior information exists as expert knowledge of an appliance's model of operation (e.g. its power demand and usage cycle), and can be encoded as a generic appliance model. This information is important as it can be used to automate the process of labelling disaggregated appliances and even be used to identify appliances which unsupervised methods cannot. It is therefore necessary to design training methods that are able to utilise both generic appliance models and aggregate consumption data without requiring sub-metered training data or complete knowledge about the type and number of all the appliances within the home.

2. NIALM TRAINING USING PRIOR MODELS AND AGGREGATE DATA

We have addressed the issues raised in the previous section in [9], particularly the problem of model training from aggregate data when sub-metered data is unavailable. To do so, we adopt a probabilistic graphical model which incorporates the difference hidden Markov model (HMM) [7], to disaggregate single appliances from household aggregate power readings. In contrast to the unsupervised training method used in [7], our approach uses generic appliance models and aggregate consumption data to generate models of specific appliance instances using expectation-maximisation (EM). Our approach then uses these trained models to disaggregate individual appliances using an extension of the Viterbi algorithm. We focus on disaggregating common appliance

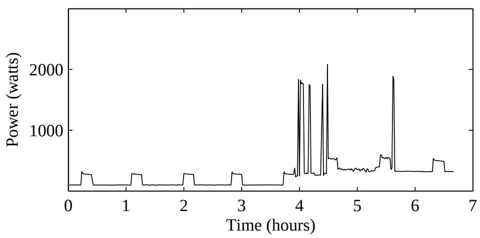

Figure 1: Example of aggregate power demand

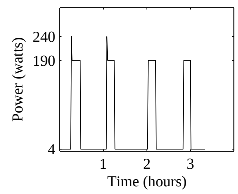

(a) Prior of appliance power.

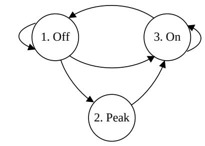

(b) State transition model.

Figure 2: Refrigerator model parameters

types which consume a large proportion of the home's energy, particularly those whose use may be deferred by the household occupants (e.g. washing machine, clothes dryer). We evaluate the accuracy of our proposed approach using the Reference Energy Disaggregation Dataset (REDD) [8].

We represent each appliance along with the household aggregate as a variant of the difference HMM. In a difference HMM, the hidden states correspond to the operational state of an appliance, while the observations correspond to the changes in aggregate power as a result of an appliance transitioning from one state to another. The difference HMM is well-suited to NIALM as it explicitly represents step changes in the aggregate power as observed data. However, since the difference HMM only considers changes in the aggregate load, it does not constrain the model such that appliances can only draw less power than the household aggregate power demand. We include such a constraint by extending the difference HMM to explicitly represent the aggregate power demand as an additional observation sequence. Both the household aggregate and difference household aggregate observation sequences complete the graphical model, and therefore training corresponds to learning the model's parameters, while disaggregation corresponds to inferring the model's hidden variables.

Our approach uses a novel training process in which prior knowledge of the generic appliance types are tuned to specific appliance instances using only aggregate data from the home in which disaggregation is being performed. The training method identifies and exploits periods during which a single appliance turns on and off without any other appliances changing state. Figure 1 shows an example of such a signature from hours 0 to 3 when only the refrigerator is turning on and off. The approach extracts such signatures by applying the EM algorithm to small overlapping windows of aggregate data. The EM algorithm is initialised with the prior appliance models and therefore restricts the behaviour the model can represent. The generic model of an appliance type consists of priors over each parameter of an appliance's model. The prior state transition matrix consists of a matrix, in which possible transitions between states are represented by a probability between zero and one, and transitions which are not possible in practice are represented by a probability of zero. The prior over an appliance's emission function consists of expected values of the Gaussian distribution's mean and variance. Figure 2 (a) shows how an expert might expect a refrigerator to operate, while (b) shows a corresponding state transition model. Once prior models have been used to extract appliance signatures from the aggregate load, the signatures are used to tune the generic models of appliance types to the household's specific appliances using a single application of EM over all extracted signatures.

We present a novel iterative disaggregation method that models each appliance's load using our graphical model and disaggregates them from the aggregate power demand. Our disaggregation method uses an extension of the Viterbi algorithm [10] which filters the aggregate signal such that interference from other appliances is ignored. The disaggregated appliance's load is then subtracted from the aggregate load. This process is repeated until all appliances for which general models are available have been disaggregated from the aggregate load.

3. EVALUATION USING REDD

We have evaluated the accuracy of our proposed approach using the REDD dataset (http://redd.csail.mit.edu/). We benchmark against two variants of our training method, where the prior is not tuned using aggregate data and when sub-metered data is used to tune the prior. We show that the disaggregation performance when using our training approach is comparable to when sub-metered training data is used [9].

To demonstrate that our proposed NIALM method is applicable in real scenarios, we applied the approach described in this paper to smart meter data collected from 6 UK households. Our system is able to use prior appliance models to identify typical behaviour patterns, such as the shower used on 'high' rather than 'eco'. Such recommendations form the basis for automatic feedback, with the potential savings provided in terms of energy, financial cost or carbon emission equivalent.

4. CONCLUSIONS

In this paper, we have summarised a novel algorithm for training a NIALM system, in which generic models of appliance types can be tuned to specific appliance instances using only aggregate data. We have shown that when combined with a suitable inference mechanism, the models can disaggregate the energy consumption of individual appliances from a household's aggregate load. Through evaluation using real data from multiple households, we have shown that it is possible to generalise between similar appliances in different households. We evaluated the accuracy of our approach using the REDD data set, and have shown that the disaggregation performance when using our training approach is comparable to when sub-metered training data is used. Future work will look at extending the graphical model to include additional information such as time of day and correlation between appliance use. In such a model, the same process of prior training as described in this paper can be applied.

5. REFERENCES

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

[2] Department of Energy & Climate Change. A consultation on smart metering for electricity and gas. Technical report, UK, 2009.

[3] Department of Energy & Climate Change. Industry's Draft Technical Specifications. Technical report, UK, 2011.

[4] G. W. Hart. Nonintrusive appliance load monitoring. *Proceedings of the IEEE*, 80(12):1870–1891, 1992.

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

[6] J. Z. Kolter, S. Batra, and A. Y. Ng. Energy Disaggregation via Discriminative Sparse Coding. In *Proceedings of the 24th Annual Conference on Neural Information Processing Systems*, pages 1153–1161, Vancouver, BC, Canada, 2010.

[7] J. Z. Kolter and T. Jaakkola. Approximate Inference in Additive Factorial HMMs with Application to Energy Disaggregation. In *International Conference on Artificial Intelligence and Statistics*, pages 1472–1482, La Palma, Canary Islands, 2012.

[8] 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.

[9] O. Parson, S. Ghosh, M. Weal, and A. Rogers. Non-intrusive Load Monitoring using Prior Models of General Appliance Types. In *26th AAAI Conference on Artificial Intelligence*, In press, 2012.

[10] A. Viterbi. Error bounds for convolutional codes and an asymptotically optimum decoding algorithm. *IEEE Transactions on Information Theory*, 13(2):260–269, 1967.

[11] M. Zeifman and K. Roth. Nonintrusive appliance load monitoring: Review and outlook. *IEEE Transactions on Consumer Electronics*, 57(1):76–84, 2011.