Advanced Engineering Informatics 32 (2017) 152–162

Contents lists available at ScienceDirect

Advanced Engineering Informatics

journal homepage: www.elsevier.com/locate/aei

Full length article

An error correction framework for sequences resulting from known state-transition models in Non-Intrusive Load Monitoring

Suman Giri, Mario Bergés *

Department of Civil and Environmental Engineering, Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh, PA, USA

ARTICLE INFO

Article history:
Received 29 August 2016
Received in revised form 17 December 2016
Accepted 18 January 2017

Keywords:
NILM
Non-Intrusive Load Monitoring
Energy management
Energy efficiency
Energy estimation

ABSTRACT

Non-Intrusive Load Monitoring (NILM), the set of techniques used for disaggregating total electricity consumption in a building into its constituent electrical loads, has recently received renewed interest in the research community, partly due to the roll-out of smart metering technology worldwide. Event-based NILM approaches (i.e., those that are based on first segmenting the power time-series and associating each segment with the operation of electrical appliances) are a commonly implemented solution but are prone to the propagation of errors through the data processing pipeline. Thus, during energy estimation (the final step in the process), many corrections need to be made to account for errors incurred during segmentation, feature extraction and classification (the other steps typically present in event-based approaches). A robust framework for energy estimation should use the labels from classification to (1) model the different state transitions that can occur in an appliance; (2) account for any misclassifications by correcting event labels that violate the extracted model; and (3) accurately estimate the energy consumed by that appliance over a period of time. In this paper, we address the second problem by proposing an error-correcting algorithm which looks at sequences generated by Finite State Machines (FSMs) and corrects for errors in the sequence; errors are defined as state transitions that violate the said FSM. We evaluate our framework on simulated data and find that it improves energy estimation errors. We further test it on data from 43 appliances collected from 19 houses and find that the framework significantly improves errors in energy estimates when compared to the case with no correction in 19 appliances, leaves 17 appliances unchanged, and has a slightly negative impact on 6 appliances.

1. Introduction

1.1. Non-Intrusive Load Monitoring and energy estimation

Non-Intrusive Load Monitoring (NILM) is a set of techniques for algorithmically inferring appliance-level energy consumption information using single point sensing of the aggregate energy consumption, typically done at the main electrical panel in a building. Interested readers are encouraged to read [1,2] for an overview on the different approaches that have been proposed to tackle the NILM problem. In general, though, there are two main types of approaches that have been suggested in the literature: those that rely on segmenting the aggregate signal according to steady-state changes (event-based approaches), and those that directly model the aggregate signal as a superposition of many source signals corresponding to the appliances (eventless approaches). Though this paper focuses solely on the event-based approaches, it is worth pointing out that there is a growing body of literature on eventless approaches, where latent variable models, such as Factorial Hidden Markov Models (FHMM), are typically employed [3,4]. Hybrid approaches that incorporate elements of both have also been developed (e.g., [5]). Online proceedings of the International and European Workshops on Non-Intrusive Load Monitoring offer a more up-to-date discussion on the latest developments in NILM [6,7].

The *event*-based[1] approach to NILM typically involves the following stages: event detection, feature extraction, classification, and energy estimation. Example algorithms that fall under this category, including some by the authors, are [8–12]. As stated earlier, the overall goal of NILM is to decompose aggregate measurements of power in a building into the power consumed by each of the appliances that are operating during the monitoring period. Given this, after detecting events, extracting their features and classifying them, the final

* Corresponding author.
E-mail addresses: sumangiri76@gmail.com (S. Giri), marioberges@cmu.edu (M. Bergés).

[1] Events are significant changes in the aggregate data being monitored (usually active power) that correspond to an appliance changing its state of operation.

http://dx.doi.org/10.1016/j.aei.2017.01.006
1474-0346/

step (energy estimation) produces estimates of the per-appliance power usage, thereby decomposing the aggregate measurements of power for the building.

Briefly, the process runs as follows: first, a signal $S[t]$ (typically active power) is monitored at the aggregate level. Then an event detection algorithm (ED) searches for change-points in $S[t]$, which are assumed to indicate when appliances change their state of operation (this also assumes that each operational state has a relatively stable power consumption during steady state). Following this, a feature-extraction algorithm (FE) extracts features associated with these events. A model is created in advance – through a training phase – to learn a function (ϕ) that maps features (X) of the events to labels (Y) corresponding to the appliances (and, optionally, the state transitions) that were responsible for these events. In the classification step, the function ϕ – learnt during training – classifies the extracted features into one of the appliance labels. Finally, in the energy estimation step, the classified event labels are used to estimate the energy consumed by each appliance (E).

Energy estimation in supervised NILM is the final step that utilizes the event labels learnt from the classification step and estimates the energy consumed by appliances with the corresponding labels. For example, let $T' = \{t'_1, t'_2, \dots t'_k\}$ represent the set of event time-stamps in $S[t]$, that results as output from the ED algorithm (here, $k \in \mathbb{N}$). Let $T = \{t_1, t_2, \dots, t_m\}$, such that $T \subset T'$ and $m \in \mathbb{N}$, represent the set of events classified as belonging to appliance $y_i \in Y$ after the FE step and subsequent classification using the function ϕ. Energy estimation for y_i, then, involves learning the changes in active power that occur during each of the time-stamps in T. These are also referred to as power deltas, and are typically extracted by calculating the difference in aggregate active power (denoted by ΔP_j) over a certain window around the event time-stamp ($t_j \in T$). Once all ΔP_j values for y_i are calculated, an ordered sequence of power deltas results. Using this, the power trace for y_i can be constructed by assuming piece-wise constant power between each element of this sequence, along with an initial power value P_0, and then performing the following sum to obtain the power value at time-stamp t:

$$\widehat{P}[t] = P_0 + \sum_{j=1}^{|R|} \Delta P_j \tag{1}$$

where $|R|$ denotes the cardinality of set $R = \{t_j : t_j < t;\ \forall t_j \in T\}$. Energy estimation for y_i is then simply the sum of the power trace over the time of interest (say t_1 to t_m), i.e.,

$$E_i = \sum_{t_1}^{t_m} \hat{P}[t] \tag{2}$$

The energy estimates can, however, be affected by errors resulting from different algorithmic steps in NILM. In [13], the authors presented a method to build Finite State Machine (FSM) models for appliances based on their power delta sequence. The FSM model is a directed graphical representation with the nodes representing the different states that can occur within an appliance and the edges representing the possible transitions between them. Fig. 1, taken from [13] shows the FSM models for a Fridge, TV, and Laptop from a publicly available dataset called BLUED [14]. Using the FSM model, it is possible to constrain the power trace and energy estimates resulting from power delta sequences to mitigate some of the errors resulting from NILM. A robust framework for energy estimation should use the labels from classification to (1) model the different state transitions that can occur in an appliance; (2) account for any misclassifications by correcting event labels that violate the extracted model; and (3) accurately estimate the energy consumed by that appliance over a period of time. In this paper, we extend an error-correction step in the framework presented in [13] by introducing an improved algorithm which looks at sequences generated by Finite State Machines (FSMs) and corrects for errors in the sequence. Here, errors are defined as state transitions that violate the FSM.

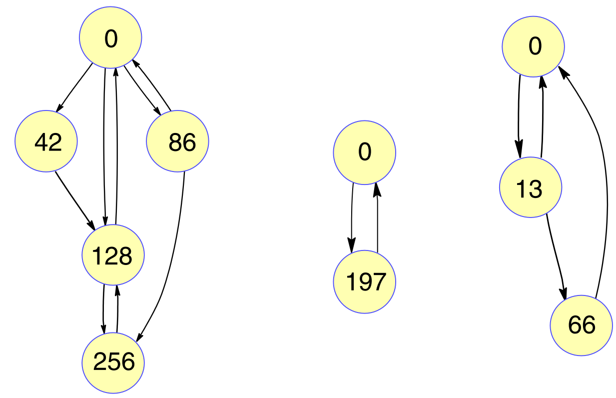

Fig. 1. FSM diagram for a Refrigerator from BLUED-1 (left), TV from BLUED-1 (center) and Laptop from BLUED-2 (right). The numbers denote the power consumption (as measured in Watts) in the different states of operation that the appliances can be in, and the arrows denote possible transitions between these states.

1.2. The need for error correction

The output of the classification step in supervised NILM is a time series of state-transition labels assigned to the observed events. It is trivial to extract ordered sequences of such labels pertaining to each appliance from this time series. Each such sequence will be denoted as C_{seq}. We assume that an operating model for the appliance that represents each such sequence is also available in the form of an FSM. Typically, this means that information about possible states, and state transitions for the appliance is available. Authors in [13] present a method to extract such information from data. Fig. 2 presents a graphical summary of the steps involved in the framework, and how error correction fits in the bigger picture. The sequence C_{seq} is a direct product of the event detection and classification steps and errors occurring in these steps (e.g., missed events, misclassifications, etc.). This results in a sequence that violates the FSM that generates it, and causes erroneous energy estimation results. Hence, an additional step of error identification and correction is required before energy estimation is performed.

The problem of correcting errors in sequences generated by FSMs has been studied extensively for problems in communication theory, DNA sequencing, pattern recognition, etc. [15,16]. Hart drew similarities between the NILM problem and the problem of decoding additive signals on a Multiple-Access Channel (MAC) [17]. He cited the low signal to noise ratio, and low message rate to channel capacity ratio as some of the desirable features of this channel. Recently, authors in [18] have begun to treat the event-based disaggregation problem from the perspective of the nascent field of Graph Signal Processing (GSP) [19–21]. In this context, error correction is performed by adding a regularizer to the problem of finding the optimal graph for total graph variation minimization. The authors report that a total graph variation regularizer (i.e., one that smooths the event label sequence) alone produces solutions that are too smooth, and incorporated simulated annealing as an error-correction proxy to refine the results. We see many connections between the work presented here and GSP-based approaches to the problem, but leave it as future work to investigate these links in more detail.

Compared to other decoding problems where error correction is required, the NILM channel presents a unique set of challenges.