

Poster Abstract: Estimating Human Interactions with Electrical Appliances for Activity-based Energy Savings Recommendations

Hông-Ân Cao
Department of Computer Science
ETH Zurich, Switzerland
hong-an.cao@inf.ethz.ch

Tri Kurniawan Wijaya, Karl Aberer
Department of Computer Science
EPFL, Switzerland
{tri-kurniawan.wijaya, karl.aberer}@epfl.ch

Abstract

Since the power consumption of different electrical appliances in a household can be recorded by individual smart meters, it becomes possible to start considering in more details the interactions of the residents with those devices throughout the day. Appliances usages should not be considered as independent events, but rather as enablers for activities. In this work, we propose an automated method for determining when an electrical device is triggered solely from its power trace. Knowing when an appliance is powered on is required for identifying recurrent patterns that could later be understood as activities. Leveraging activity knowledge over time will allow us to design personalized energy efficient measures. We envision the design of future ambient intelligence systems, where the smart home can optimize the energy consumption in regards to the lifestyles of its residents.

Categories and Subject Descriptors

G.3 [Probability and Statistics]: Time series analysis;
H.2.8 [Database Applications]: Data mining

General Terms

Measurement, Algorithms

Keywords

Smart Meters, Energy Disaggregation, Activity Inference

1 Introduction

The future smart grid offers the possibility of having fine-grained information and capabilities to monitor its status in real time. Implementing real-time and personalized feedback would amount to about 12% of energy reduction in the residential segment [1]. This relatively small figure should be contrasted with the potential savings during peak time, when high penalties might become a reality in the future. Focusing on the household scale offers an alternative to aggregating levels in Demand Response Systems. In the context of

the smart home, one could foresee trading-off users' lifestyle preferences and comfort with saving measures, while preserving the privacy of the residents.

It has yet to be decided how much information should be collected, i.e., the granularity of such data, and which additional sensors should be integrated to provide a better understanding of how energy is consumed. To this end, the access to disaggregated data requires the setup of data collection architectures or the development of NILM algorithms on existing household-level aggregated data to differentiate the devices being used. Given the recent release of a large dataset with appliance-level measurements, abstracting the usage of electrical devices in households by investigating the motives behind them being triggered by a user becomes possible. This involves unraveling information from the power measurements collected and finding out when and how they are used in conjunction. This stands in contrast to prior work where the power traces were simulated [4]. Chen et al. [2] used statistical attributes of the data to determine occupancy, we are however assessing activities that incur energy consumption.

In order to determine which appliances are utilized jointly and linked to a human activity, our contribution is to distinguish the "active" consumption from the baseline and noise in their power traces. The novelty and the challenge reside in having no side information that could assess the proximity of the residents, nor ground truth from a journal that documents the activities in the household.

2 Methodology

Using only electrical loads (no side information, nor ground truth), it is necessary to evaluate how to differentiate between baseline consumption that can be considered as noise, and human-triggered actions. While it would be possible to handpick a threshold to decide when the appliance is powered on and serving a human activity, such process would be done arbitrarily and would not be generalizable across the whole population given the multitude of brands and models in consumer electronics. To this end, we developed an automated way of deciding when an appliance reaches a power level high enough, such that it can be regarded as being used by a human being. This requires considering each household separately and learning from the specificity of each trace. Such method relates to image thresholding [3].

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the Owner/Author.

Copyright is held by the owner/author(s).

BuildSys'14, November 5–6, 2014, Memphis, TN, USA.
ACM 978-1-4503-3144-9/14/11.
<http://dx.doi.org/10.1145/2674061.2675037>

2.1 State Estimation

We consider two types of power traces, namely appliance level data (single appliances), and circuit-level data (aggregated readings from room-level, or power strip data). We refer to both as *appliances* from now on. We explain how different power levels are linked to the appliance’s state and its utilization. Since a human being is not activating the appliances throughout the day, we can distinguish between an *idle* state (off/stand-by mode, typically low power levels) and an *active* state (when the residents are powering it on or actively interacting with it). We notice, for example, in the case of a washing machine, that several mechanisms allow running different washing programs and steps (soaking, spinning, etc.), while for a room, we could expect to see different devices (lights, smaller consumer electronics) being turned on. So, we rely on this to suggest that different states in the use of an appliance are linked to different levels of power. Following this idea, we want to observe the relationships between power levels in the distribution of the power measurements of an appliance.

Although we intend to discover activities in a data-driven manner, i.e., without a-priori knowledge, nor human labeling, we have in mind for the time-being high level activities (such as cooking, cleaning, etc.). This means that we do not dwell into the intricacy of the different stages involved in an activity (in the case of cooking: cleaning vegetables, heating ingredients, eating, etc.). Thus, if we consider a power strip in the kitchen and its respective power readings, the transitions in the traces might be due to smaller appliances being powered on (kettle, mixer, etc.). However, since, they are not disaggregated, they cannot be labeled and cannot be directly used. This is why we focus on the overall duration of the interaction with an appliance, not differentiating between all the stages and sub-activities it might involve.

2.2 Gaussian Mixture Model

Given that most of the appliances operate at low power levels during their idle period, this state can be identified as the first set of correlated measurements. We model the distribution of power levels by approximating it with a Gaussian Mixture Model (GMM). The threshold between on and off lies in the first valley of the Gaussian mixture (the first Gaussian represents the idle status, while starting from the second Gaussian, the appliance is in use).

3 Experimental Results

The PecanStreet dataset (<http://wiki-energy.org/>) comprises 239 households with 1-minute readings from January to May 2014. While 73 different types of readings were expected to be collected, there are at most 22 actively monitored circuits per household. Appliances with larger ranges of consumption are for example ovens or dishwashers, with power readings up to 3 kW. We exclude always-on devices (such as fridges) and those consuming less than 0.5 Wh per week.

Our algorithm considers one month of data per appliance (to minimize the impact of weather). We observe for each month that some power levels readings amount to thousands of occurrences, while the magnitude of other representatives is in the order of hundreds to a few instances. Because of

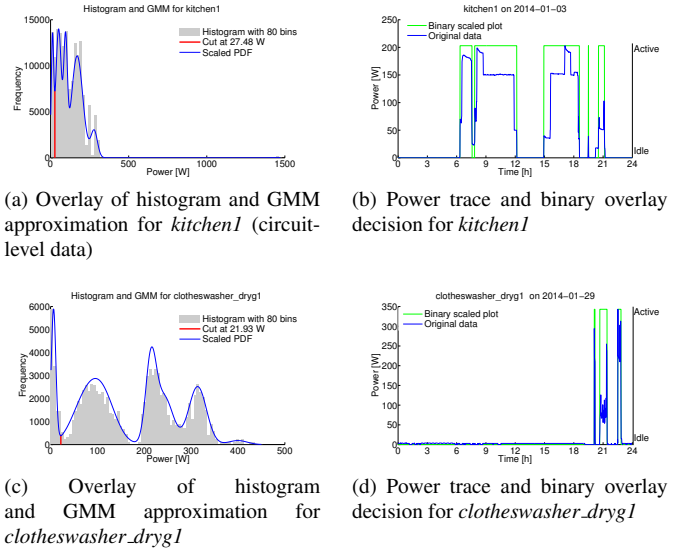


Figure 1: Outcome of the GMM. In (a) and (c), power below the threshold is considered to be in the idle state, and in the active state otherwise.

this, the data are scaled to lessen the order of magnitude between the measurements. For this purpose, we resample the data, such that for each power value i , their quantity n_i is of the order of $C * \log(n_i + 1)$, where C is a constant, which we set to 400 to guarantee that enough data are available.

We use a parametric implementation for the GMM from Matlab. We select the best model by choosing the lowest Bayesian Information Criterion (BIC) value corresponding to the number of Gaussians in the mixture. We show in Figure 1 the outcome of the GMM for a particular household, where the threshold for the active state is set at 27.48 W for *kitchen1* and 21.93 W for *clotheswasher_dryg1* respectively.

4 Future Work

In this work, we introduced an automated way of determining when an appliance is activated by a human being by filtering out baseline noise from the readings and by looking at the distribution of the power measurements. Having now obtained binary vectors of data, we intend to consider time intervals throughout the days and infer patterns of appliances being used conjointly and derive temporal rules. We will improve the evaluation of our method by comparing the results to human annotated data.

5 References

- [1] K. Carrie Armel, A. Gupta, G. Shrimali, and A. Albert. Is disaggregation the holy grail of energy efficiency? The case of electricity. *Energy Policy*, 52:213–234, 2013.
- [2] D. Chen, S. Barker, A. Subbaswamy, D. Irwin, and P. Shenoy. Non-Intrusive Occupancy Monitoring using Smart Meters. *BuildSys’13*, pages 1–8, 2013.
- [3] Z.-K. Huang and K.-W. Chau. A new image thresholding method based on Gaussian mixture model. *Applied Mathematics and Computation*, 205(2):899–907, 2008.
- [4] I. Richardson, M. Thomson, D. Infield, and C. Clifford. Domestic electricity use: A high-resolution energy demand model. *Energy and Buildings*, 42(10):1878–1887, 2010.