A Temporal Motif Mining Approach to Unsupervised Energy Disaggregation

Huijuan Shao
Dept of Computer Science
Discovery Analytics Center
Virginia Tech, VA 24061
huijuans@vt.edu

Manish Marwah
HP Labs
Palo Alto, CA, 94034
manish.marwah@hp.com

Naren Ramakrishnan
Dept of Computer Science
Discovery Analytics Center
Virginia Tech, VA 24061
naren@cs.vt.edu

ABSTRACT

Non-intrusive appliance load monitoring has emerged as a popular approach to study energy consumption patterns without instrumenting every device in an installation. The resulting computational problem is to disaggregate total energy usage into usage by specific circuits and devices, to gain insight into consumption patterns. We exploit the temporal ordering implicit in on/off events of devices to uncover motifs (episodes) corresponding to the operation of individual devices. Extracted motifs are then subjected to a sequence of constraint checks to ensure that the resulting episodes are interpretable. Our preliminary results show that motif mining is adept at distinguishing devices with multiple power levels and at disentangling the combinatorial operation of devices.

1. INTRODUCTION

As the saying goes, sustainability begins at home. Greater than ever before, there is now a significant interest in reducing household energy footprints by providing consumers with detailed feedback on their energy consumption patterns. By contrasting such 'drill-down' data with neighborhood profiles, consumers can make better informed decisions about how their daily activities impact the environment as well as their bottom line.

A key step in this endeavor is energy disaggregation. This is the task of, non-intrusively, monitoring aggregate energy usage (electricity, water) at a home/unit and separating it out into individual appliances, subunits, and other spatial dimensions, automatically using machine learning methods. A variety of methods have been proposed, e.g., factorial HMMs [3] and sparse coding [2] but the increasing diversity of appliances to be accommodated and the spatio-temporal coherence properties that must be modeled provides continuing opportunities for algorithm innovation.

Here we propose a temporal motif mining approach [5; 6] to energy disaggregation. We specifically focus on low-frequency measurements and aim to characterize stable power consumption events, in contrast to transients. The basic idea is to discover the minimal episode which corresponds to a complete state-change cycle by a device or part of a device. Unlike state-of-the-art probabilistic methods that posit detailed temporal relationships and involve complex inference steps, we argue that our method is lightweight and, at the same time, capable of accuracy levels better than or comparable to these more complex methods. Using this approach, we conduct a thorough experimental investigation of our method on the REDD [4] dataset, demonstrating the ability of our approach to scale w.r.t. the number of devices and, at the same time, achieve stability of disaggregation accuracy.

2. ALGORITHM

A complete state-change cycle of a device captures the device's operation from when it comes on to being switched off. But because we have access to only aggregated data, one complete state-change cycle will be interrupted by the operations of other devices. Motif mining aims to disentangle such seemingly discontinuous profiles. We will show how this approach has three inherent advantages over HMM-based algorithms: (i) it leverages local methods for pattern discovery; (ii) it can naturally scale to multiple devices, and multiple states for a single device; and (iii) it is robust to missing data.

A typical home usually receives 2-phase power from the utility with each phase connected to many circuits. Each circuit in turn has one or more devices that draw power from it. High power consumption devices typically connect to both the phases. Given only the aggreagte power consumption data for each phase, in order to disaggregate the consumption into individual devices requires us to build a working motif model for the operation of each device.

The key steps of our approach are outlined in Fig. 1 where the input is the aggregated power observation time series and the output is the disaggregated time series for each device. First, rather than modeling the power consumption, we model power change events (what we refer to as the 'diffs data.'). From the diffs data, we remove rapid spikes in observations, as might manifest when devices are turned on or off. The power states for such spikes typically have a range of (-10%,10%). Traditionally, the power state for a given device conforms to a Gaussian distribution rather than a fixed level.

In the second step, we exploit motif mining techniques based on frequent episode algorithms. The individual symbols of the episode are state transitions, so we are looking for frequently recurring sequences of transitions. As shown in Fig. 1, two instances of one episode: (+610, -605), (+600, -600) are discovered. Note that we allow for some tolerance in defining power levels when counting occurrences. This step typically results in many episodes that will be the subject of pruning algorithms later. Usually when (+600, -600)

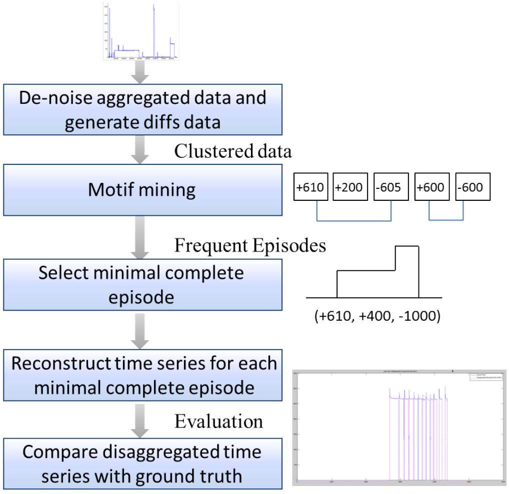

Figure 1: Disaggregation using temporal motif mining.

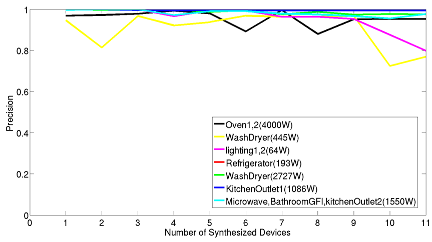

Figure 2: Precision of disaggregation as a function of number of aggregated circuits.

is frequent, (-600, +600) is also frequent.
In the third step, we aim to identify episodes which reflect key electricity consumption characteristics. Such interpretable episodes are chosen by imposing a series of constraints. One constraint states that the net sum of power levels in an episode must be near zero. That would exclude episodes such as (+600, -1000). A second constraint is that the first k transitions in the episode should sum to a positive power level. For instance, the episode (+600, -1000, +400), although satisfying the first constraint, could not possibly be realized by a single device. Finally, we select *minimal* complete episodes, with a view toward identifying those transitions that could have been generated by only a single device (or part of a device). For instance, if we are given two episodes (+600,+400,-400,+600) and (+600,-600), then (+600,+400,-400,-600) is excluded because it is not minimal. The third constraint is crucial since it helps identify unseparatable episodes and which would correspond to devices.
Finally, the extracted minimal complete episodes are reconstructed to generate time series for each putative device; in the experiments below we compare the disaggregated time series with the ground truth time series for each device.
We conduct experiements using the REDD low frequency data set from [4]. From the ground truth data, we synthesize aggregate circuit data, ranging from 1 to 18 circuits. For each of the aggregated datasets, we apply our multi-level approach and compare it with the ground truth. We primarily employ the measure of circuit level precision (the amount of energy corresponding to each circuit). Fig. 2 shows that circuit level precision is relatively stable with increase in synthesized circuits from 1 to 11. Totally eleven synthesized circuits can be disaggregated. There are 7 minimal episodes corresponding to these 11 devices. Since Oven(4000W) connect to two circuits, therefore, only one device for two circuits is identified. And two lighting1,2(64W) share similar power consumption, therefore, these two lights cannot be distinguished. We combine the ground truth of these two lights then compare with disaggregated episode(+64W,-64W). We found that the less frequency of episode generated by device,such as WashDryer(445W), or if the power consumption is low, e.g. lighting1,2(64W) the precision is prone to fluctuate with the increase of synthesized devices.

3. DISCUSSION

We have described an intuitive motif-based approach to disaggreation that performs well relative to more complex algorithms that perform detailed modeling of temporal profiles. More importantly, we have demonstrated how our approach is not just an aid to disaggregation but, as a byproduct, also extracts temporal episodic relationships that shed insight into consumption patterns. In this sense, our work goes further than past work into addressing the real goal of disaggregation research, viz. to understand systematic trends in consumption patterns with a view toward identifying opportunities for savings.

4. REFERENCES

[1] G.Hart. Nonintrusive appliance load monitoring. Proceedings of the IEEE 1992: 80(12).

[2] J. Zico Kolter, Siddharth Batra, Andrew Y. Ng: Energy Disaggregation via Discriminative Sparse Coding. NIPS 2010:1153-1161.

[3] Hyungsul Kim, Manish Marwah, Martin F. Arlitt, Geoff Lyon, Jiawei Han: Unsupervised Disaggregation of Low Frequency Power Measurements. SDM 2011:747-758.

[4] J.Z.Kolter, Matthew Johnson. REDD: A Public Data Set for Energy Disaggregation Research. SustKDD 2011.

[5] Bill Yuan-chi Chiu, Eamonn J. Keogh, Stefano Lonardi: Probabilistic discovery of time series motifs. KDD 2003: 493-498.

[6] Dragomir Yankov, Eamonn J. Keogh, Jose Medina, Bill Yuan-chi Chiu, Victor B. Zordan: Detecting time series motifs under uniform scaling. KDD 2007:844-853.