Non-Intrusive Model Derivation: Automated Modeling of Residential Electrical Loads

Srinivasan Iyengar, David Irwin, and Prashant Shenoy
University of Massachusetts Amherst

ABSTRACT

A variety of energy management and analytics techniques rely on models of the power usage of a device over time. Unfortunately, the models employed by these techniques are often highly simplistic, such as modeling devices as simply being on with a fixed power usage or off and consuming little power. As we show, even the power usage of relatively simple devices exhibits much more complexity than a simple on and off state. To address the problem, we present a Non-Intrusive Model Derivation (NIMD) algorithm to automate modeling of residential electric loads using concepts from power systems, statistics, and machine learning. NIMD automatically derives a compact representation of the time-varying power usage of any residential electrical load, including both the device's energy usage and its pattern of usage over time. Such models are useful for a variety of analytics techniques, such as Non-Intrusive Load Monitoring, that have relied on simple on-off models in the past. We evaluate the accuracy of our models by comparing them with both actual ground truth data, and against models that have been designed manually by human experts. We show that models derived via NIMD are comparable in accuracy to models built by experts and closely approximate the ground truth data.

CCS Concepts

•**Computing methodologies → Model development and analysis; Model verification and validation;**

Keywords

Energy metering, Load modeling and analysis

1. INTRODUCTION

Buildings constitute nearly 40% of the total energy and 75% of the total electricity consumption in advanced economies, exceeding the usage from other large sectors, such as manufacturing and transportation [14]. Since homes and residential buildings comprise nearly half of this usage, techniques to reduce the energy footprint of residential buildings has received significant attention in recent years. Most of these techniques rely on a clear understanding of what electrical loads are present in a home, how much electricity is consumed by each load, and how electricity is used by residents (e.g., their daily activities). An understanding of these factors is central to a broad swath of energy management and optimization techniques, including *automated demand-response*, where certain loads are turned off at the request of the grid, *automated load control*, where loads are automatically turned on or off to match user activity, and *energy analytics*, where algorithms analyze energy usage data to derive and exploit behavioral insights to implement various optimizations. Such techniques are becoming more commonplace with the growing popularity of Internet-of-Things (IoT) devices being deployed in smart homes.

We argue that modeling of electrical loads is a fundamental building block that is essential for driving higher-level techniques such as automated DR or energy analytics. Such models are compact representations of the electrical usage patterns exhibited by a load (e.g., a washing machine or TV) as well as temporal characteristics that describe when the load is used by residents (e.g., residents watch TV every evening and do laundry on weekends). Studies have observed that simplistic or coarse-grain models can be detrimental to the accuracy and effectiveness of higher-level approaches. For instance, Non-Intrusive Load Monitoring (NILM) is a well-studied energy analytics technique that disaggregates overall energy consumption data for a home into individual loads [12, 16, 1]. NILM-based analytics approaches have been used in a variety of applications, such as inferring occupancy patterns [9, 15], reducing peak demand by opportunistic load scheduling [6], and learning thermostat schedules [13]. However, many NILM approaches still model electrical loads as simple *on-off* devices, where the load draws a fixed amount of power when turned on. As illustrated in Figure 1, which depicts a washing machine, most residential loads exhibit complex and varied power patterns that are distinct from the simple *on-off* behavior. Hence, disaggregating loads based on a simplistic and inaccurate understanding of a load's behavior significantly degrades the accuracy of higher-level techniques.

Despite its importance, empirical or analytic modeling of electrical loads has received relatively little attention. Typically, common devices, such as TVs, refrigerators, and computers, are only rated (often conservatively) based on their maximum power, but include no details of how the device consumes power over time. A recent effort [3, 4] analyzed empirical data gathered from a large number of residential loads to argue that only the simplest loads, such as light bulbs, exhibit a simple on-off behavior and demonstrates that most loads exhibit more complex exponential decays or growth, bounded min-max, and cyclic patterns. While this work proposed more complex analytic models to describe load behavior, it did not propose any algorithms or approaches to derive (or construct) such models automatically. That is, it required manual modeling of a load by an expert before such a model could be used by higher-level optimizations. However, manual modeling of highly complex loads is time-consuming, and, for a load as complex as Figure 1, potentially infeasible given the load's complexity.

Thus, in this paper, we propose an automated modeling approach for residential electrical loads. Our approach, which we refer to as

e-Energy'16, June 21 - 24, 2016, Waterloo, ON, Canada

ISBN 978-1-4503-4393-0/16/06…$15.00

DOI: http://dx.doi.org/10.1145/2934328.2934330

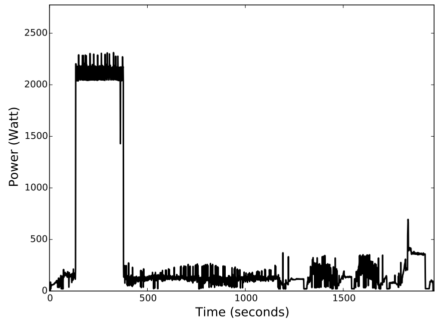

Figure 1: Observed power usage of a washing machine

Non-Intrusive Model Derivation (NIMD), takes a trace of power usage for an electrical load and automatically constructs an analytical model that captures both the device power characteristics and a temporal usage model of the load. Our approach is fully automated and does not require any manual intervention making it suitable for a range of higher-level energy management and optimization techniques. Our goal is to enable simple construction of highly detailed power models for any device. We believe that such models could be used for a wide range of applications. Given NIMD, users could easily construct and post detailed models of device energy usage, similar to the simple models provided by the crowd-sourced Power Consumption Database[1], which currently only records a static normal and standby power for devices but includes no details of how they consume power over time. Alternatively, NIMD could provide device manufacturers a methodology for modeling their own devices at manufacturing time, enabling them to release more detailed power usage models than simple maximum power consumption readings. In designing our Non-Intrusive Model Derivation approach, we make the following contributions.

Modeling Challenges. We review common features in the power usage of different types of loads based on their electrical characteristics from prior work. We then highlight the challenges in automatically modeling complex loads based their power usage data. Due to these challenges, prior work assumes that human experts identify the model type and manually construct it, even though such manual modeling is both time-consuming and error-prone.

NIMD Algorithm. We present an approach to automatically derive a load model solely from opaque time-series power data. Our NIMD algorithm leverages a set of techniques from power systems, statistics and machine learning to first derive a device usage model that extracts load active periods, detects complex state changes and cycles in the load operation, and parameterizes and fits basic models onto loads to derive a compact representation of a device's power usage profile. Our algorithm then derives a usage model that represents when and how often the device is used.

Model Evaluation. We evaluate the accuracy of our models by comparing them with actual ground truth data, and against models that have been designed manually by human experts. We show that models derived via NIMD are comparable in accuracy to manually designed models and closely approximate the ground truth.

The rest of this paper is structured as follows. Sections 2 and 3 present the problem statement and background on electrical loads. Section 4 presents our NIMD approach. Sections 5 and 6 presents our implementation and experimental results. We present related work in Section 6 and conclude in Section 7.

2. PROBLEM STATEMENT

A typical home or residential building has dozens of electrical loads of various diverse types and sizes. Some common loads include lights, HVAC equipment (such as AC and heaters), appliances (such as a washing machine and dishwasher), and electronic equipment (such as the TV, music system, phones, and chargers). Different loads will exhibit different usage patterns. A few loads such as a clock may be always on, but most loads are active only when they are in use by their users and are off or on standby at other times. When a load is active, it draws a certain amount of power, which may vary over time based on its operation. Similarly, when a load is inactive or in standby mode, it will draw either no power or a small amount of standby power. In general, a load may draw both real and reactive power, each with a distinct power pattern. Real and reactive power derive in an AC system when the voltage and current phase are not precisely aligned. As a result, at some moments, the product of current and voltage will be negative, causing power to flow back towards the generator. The portion of productive power that flows toward the load is real power, while the portion of power that flows back toward the generator is reactive power. For simplicity, this paper focuses only on modeling real power. However, our basic approach is easily extended to modeling reactive power, which we leave for future work.

We assume that a load can be empirically monitored—for example, using commonly available outlet sensors such as a Belkin Wemo [22] or Kill-a-watt—which yields a trace of its power usage over time. The goal of our work is to *automatically* derive a compact model for the load based on its power data that captures and describes its power usage over time. Intuitively, the raw power trace of a load is itself a model, since it fully describes the load's behavior over time; however a raw trace is voluminous and not useful as an input for higher-level energy optimizations. Thus, our goal is to develop models that are compact analytic descriptions of a load's behavior. Any compact model will necessarily be an approximation of the load's raw power trace and there may be a tradeoff between the accuracy of the model and the compactness of its representation. Given a load's raw power trace, our goal is to automatically derive a model that appropriately balances these tradeoffs.

Formally, a model for an electrical load has two key components: (i) a *device model* that captures how the load behaves when active, and (ii) a *usage model* that captures when and how frequently a load is used. Note that the device model is inherent to the device and its characteristics, while the usage model is inherent to the how users use that load. In other words, a device model for a certain load is an invariant across homes, while its usage model may vary from home to home. For example, a certain model of a washing machine will exhibit the same power usage characteristics when turned on (with the same device model), but different users may choose to do their laundry at different times and frequencies—daily or on weekends–yielding different usage models.

Thus, given a raw power trace of a load $X = (X_1, X_2, \ldots X_k)$ over k time units, where X_i denotes the (real or reactive) power used by the load at time i, we seek to automatically learn (i) a device model comprising a set of functions $\langle f^{([1..n])} \rangle$ that describe its power usage over time when active, and (ii) a usage model that comprise a set of probability distribution functions $\langle g^{([1..m])} \rangle$ that describe when and how frequently the load is activated.

3. CHARACTERIZING LOADS

In this section, we summarize the various kinds of electrical loads found in residential environments. This understanding is critical for automated modeling of such loads. We also describe how various kinds of electrical loads that are typically used by their users. Our characterization builds on prior work on empirically characterizing electrical loads [4, 3].

[1]http://www.tpcdb.com/

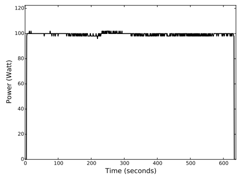

(a) Lamp (Resistive)

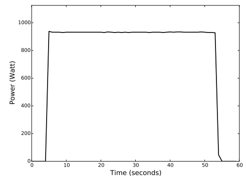

(b) CoffeeMaker (Resistive)

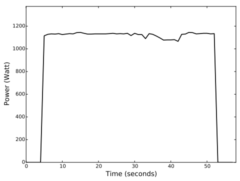

(c) Vacuum Cleaner (Inductive)

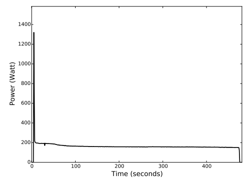

(d) Freezer (Inductive)

Figure 2: Examples of Resistive and Inductive loads

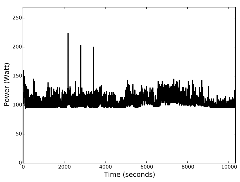

(a) PC-Desktop (Non-linear)

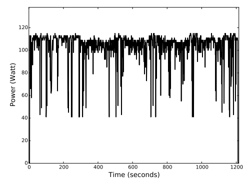

(b) LCD-TV (Non-linear)

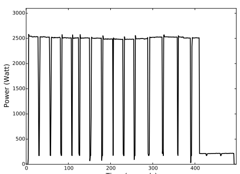

(c) Laundry Dryer (Composite)

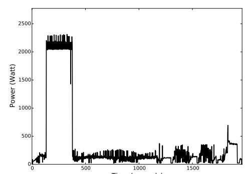

(d) Washingmachine (Composite)

Figure 3: Examples of Non-linear and Composite loads

3.1 Common Residential Load Types

Elementary power systems divides electrical loads into four basic types: **resistive**, **inductive**, **capacitive** and **non-linear**.

Resistive Loads: A pure resistive load is any device with a heating element where electrical energy is dissipated as heat or light. Example residential devices that fall into this category include incandescent lights, toasters, ovens, electric stoves, space heaters, and electric water heater (Figure 2). A resistive load draws power such that its voltage and current waveforms are aligned and in phase.

Inductive loads: An inductive load is one where the current waveform lags the voltage waveform. The most common type of household inductive load is any device with a motor, such as fans, food processors, vacuum cleaners, and any load with a compressor, such as a refrigerator or air conditioner (see Figure 2). Typically, inductive loads draw an initially high amount of power, followed by a lower steady power draw [4]. The initial power "spike" is caused by the higher in-rush current drawn by a motor when it starts up, followed by a steady current to run the motor at a steady speed.

Capacitive loads: A capacitive load is one where the voltage waveform lags the current waveform. An example capacitive load is an "overexcited" synchronous motor that has significant current in its rotor windings. Capacitive loads tend to be common inside the core of the electric grid, e.g., long transmission lines, capacitor banks for power factor correction, but they are largely absent in residential environments. Given our focus on residential loads, we ignore capacitive loads in the rest of this paper.

Non-linear loads: A non-linear load is one where the waveform of the current does not follow the sinusoidal waveform of the applied voltage. This is caused by a switching action of the load, resulting in a non-sinusoidal current draw. Common electronic devices use switched-mode power supplies that can continuously vary the current draw of the device to meet its needs, resulting in non-linear loads. Such loads include any electronic device with a power supply, such as a TV, music system, computer, phones, chargers of various sorts; these power supplies are capable of supplying a variable amount of current to the load (see Figure 3).

Composite loads: While many simple electrical loads fall into one of the above three categories (resistive, inductive, non-linear), many household devices comprise one or more of these base load types, resulting in a more complex compound load. For instance, a refrigerator has a compressor (an inductive load) and a light bulb (a resistive load) that turns on when the door is open. A dishwasher has a motor (an inductive load), a heating element (a resistive load) to heat water, and a pump (also an inductive load) to drain water. Such loads may activate their individual component loads in sequence or parallel and the resulting power draw is the sum of the power drawn by each component load (see Figure 3).

3.2 Characterizing Usage

Residential loads can be classified as a background loads and foreground loads. A background load is one that typically runs in the background without active user intervention. A refrigerator is the most common background load that is always on and controlled by an internal thermostat that periodically turns on the compressor. Typical background loads exhibit periodic behavior, where they turn on for some duration and then become inactive until the next period. Thermostat-controlled equipment such as a central AC, furnace, tank-based water heaters, are all examples of background loads that also exhibit periodic usage behavior. The period may vary over time due to changes in the environment i.e. longer period of an AC due to higher outside temperature.

Foreground loads are ones that are actively controlled by a user—the user turns them on when needed and turns them off when done. The usage of foreground loads depends on the type of the load and how users operate them in their daily routines. Each load will have a usage frequency that governs how frequently it is turned on by the user (e.g., daily, weekly) and the usage may have time-of-day, day-of-the-week or seasonal patterns associated with them. For example, lights may be used in the evenings when it is dark; laundry may typically be done on weekends, heaters may only be used in the winter. The frequency of use and temporal characteristics govern the usage patterns of both foreground and background loads.

3.3 Manual modeling of loads

A recent effort [3, 4] empirically characterized electrical loads from their power traces and demonstrated that common loads can be modeled analytically. In particular, the study showed that the empirically observed behavior of each basic load type can be modeled using one of four analytic equations:

On-Off Model: In this case, the load draws a fixed power P_{on} when active and zero or a small amount of standby power P_{off} when inactive. Simple resistive loads were found to exhibit such binary on-off behavior. Figures 2(a) and (b) shows examples of the power usage of resistive loads with on-off behavior.

On-Off Decay Model: In this case, the power usage of the load exhibits an exponential decay behavior, represented as follows.

$$\mathbf{X}(t) = \begin{cases} p_{active} + (p_{peak} - p_{active})e^{-\lambda t}, & 0 \le t < t_{active} \\ X_{off}, & t \ge t_{active} \end{cases} \quad (1)$$

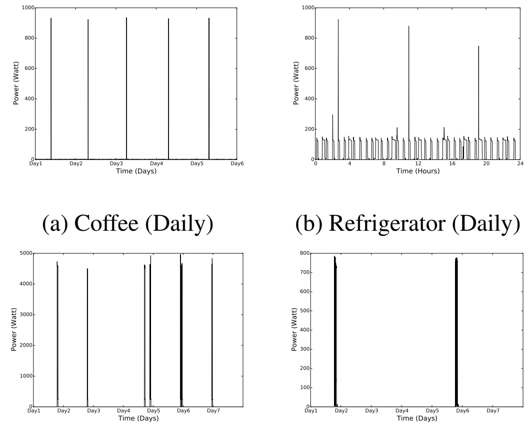

(a) Coffee (Daily) (b) Refrigerator (Daily)

(c) Clothes Dryer (Weekly) (d) Dishwasher (Weekly)

Figure 4: Examples of loads with Usage characteristics

Here, p_{peak} represents the initial surge power, p_{active} is the stable power level and λ captures the rate of decay. Many inductive loads consisting of AC motors were shown to exhibit this behavior. Higher-powered resistive loads were also shown to exhibit an exponential decay, but with a less prominent peak and a gentle decay. Figure 2(c) and (d) shows examples of the power usage of inductive loads with a spike and exponential decay behavior.

On-Off Growth Model: Some loads exhibit a growth behavior i.e. a logarithmic growth in power usage. We model such devices using a logarithmic function (inverse of the exponential function) that starts with a power level p_{base} with a growth parameter λ. We refer to such loads as an *on-off growth* model:

$$\mathbf{X}(t) = \begin{cases} p_{base} + \lambda \cdot \ln t, & 0 < t < t_{active} \\ X_{off}, & t \geq t_{active} \end{cases} \tag{2}$$

Stable Min-Max and *Random Range* Models: All non-linear loads exhibit a degree of random behavior and the observed behavior was characterized as a random walk between an upper and lower bound (referred to as random range) or a stable power draw with random upward or downward deviation (referred to as a stable min-max model). The random variations were modeled using a uniform distribution. Electronic loads with switched-mode power supplies, such as TVs, phone chargers, and computers were shown to exhibit this behavior. Figures 3(a) and (b) shows two examples of non-linear loads with random variations.

Cyclic Model: Any load that exhibits repeating patterns was characterized as cyclic with a certain period. All other complex loads that included multiple types of basic loads were characterized as a linear combination of above loads. Figures 3(c) and (d) show examples of cyclic loads that repeat and a composite load that embeds the operation of multiple basic loads.

The previous study [3, 4] involved manual modeling of electrical loads by an expert. It assumed that the load type (e.g., TV, washing machine, AC) was known *a priori* and that a human expert could map these loads onto one of the above analytic models based on their knowledge of whether the load was resistive, inductive, or non-linear. The parameters of the analytic model chosen for the load were then manually derived. This human expertise was particularly important to model complex loads such as a washing machine (see Figure 1) where human expertise is used to model each type of power variation observed in the trace. Thus, at best, this method lends itself to a supervised approach where a human expert uses her domain knowledge of the type of device to label the type of model to be used for the load, and parameters of the model are then manually derived using the empirical trace data.

The problem addressed in this work is more challenging. We assume that our system is provided with a power trace with *no a priori information* of what type of device it was gathered from or any knowledge of the load type. We seek to automatically (i.e., unsupervised) derive the "best" analytic description that explains the observed behavior. Further, we seek to automatically derive both the device model and the usage model, while prior work [3, 4] only dealt with device models and not how the device was used in a particular environment. As discussed earlier, such analytic load models are potentially useful in a wide range of scenarios.

Extensions: While prior work [4] employed a uniform distribution to model random power fluctuations in non-linear loads, our work uses the more general $Gamma$ and $LogGamma$ distributions to model stable min and stable max behavior with random deviations. The two models are shown below.

$$\mathbf{X}(t) \sim \; Gamma(\alpha, loc, scale), \quad 0 < t < t_{active} \tag{3}$$

$$\mathbf{X}(t) \sim \; LogGamma(\alpha, loc, scale), \quad 0 < t < t_{active} \tag{4}$$

Here, α, loc and $scale$ denote the shape, location and scale parameters for the two distributions. Further, for random range devices a normal distribution could be used. Apart from these, more load types can be used with future appliances having varied characteristics as they get added to residential buildings.

4. NIMD ALGORITHM

In this section, we propose our Non-Intrusive Model Derivation (NIMD) approach for automated (unsupervised) modeling of electrical loads. Broadly our approach has two parts: (i) device modeling, where we learn the power usage behavior of the load when it is active, and (ii) usage modeling, where we learn how the users use the load in a particular environment. Although, both components are necessary to model the overall load behavior, they are independent and can be used on their own for specific use-cases. In what follows, we first describe our basic approach, followed by the details of the device and usage modeling.

4.1 Basic Approach

Figure 5 depicts the high-level approach for NIMD device modeling. Given a raw power trace of a load, NIMD's approach to constructing a device model involves the following steps:

Step 1. *Active period extraction:* For a given trace, the first step is to partition the trace into active and inactive periods. An *active* period is one where the load is operating and drawing power, while an *inactive* period is one where the load is turned off or in standby mode (and not in active use). A long power trace will consist of alternating periods of active and inactive use, and hence, this step extracts active periods from the trace.

Step 2. *State change detection via change point detection:* During each active period, a load may transition through different active states and exhibit a different type of power variations in each state as it transitions from one active state to another. In this step, our technique uses a change detection algorithm to determine these state transitions, which manifest as "significant" changes in power behavior. By further partitioning an active period at each state transition, we obtain a set of trace segments corresponding to different active states within each active period.

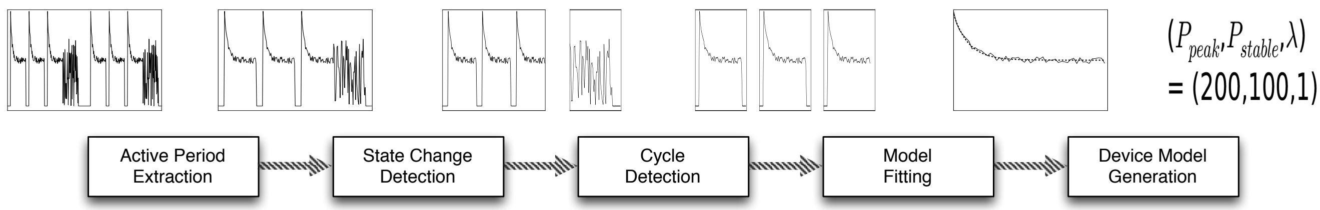

Figure 5: Pipeline of steps involved in devce modeling of an electrical load

Step 3: *Cycle detection:* Next our technique compares the power patterns across states to determine if the behavior is cyclic. If a repeating pattern of state transitions is found, then a more compact model can be constructed by analyzing a repeating cycle rather than all trace segments.

Step 4: *Model fitting:* In this key step, our technique tries to fit the trace segment extracted from each state onto various analytic models described in Section 3. The best fit is then chosen, which yields both the load type seen during that active state and the parameters of the model describing that observed behavior.

Step 5: *Device Model Generation:* The previous step yields a sequence of analytic models, one for each active phase, as well as cyclic dependencies, if any, for each active period. We repeat this process for each active period present in the trace. The final step then is to catalog the sequence of analytic functions for the overall model as well as the parameters of the various analytic functions found by our technique.

Usage Modeling: While the previous steps construct a device model from a raw power trace, we now describe the high-level approach for deriving a usage mode for the load.

Intuitively, the usage model involves determining *how frequently* a load is used and *when* it is used (e.g., mornings, evenings, weekends, summer, etc.) To derive the usage, consider the first step of the device modeling, namely active period extraction. In this step, the trace is partitioned into active and inactive periods. In doing so, we obtain, over the period of the trace, start times of each active period, and the lengths of each active ("on") and inactive ("off") periods. This data is used to construct a usage model as follows:

Step 1: Our technique first finds the shortest duration (e.g., a day, week, month or year) over which the load exhibits "similar" behavior. In order to derive a compact model, this is the period over which the usage of the device repeats in a statistically meaningful manner and captures the seasonality of the usage.

Step 2: Next, our technique constructs probability distributions for the start times and the active and inactive period lengths for the above duration. The joint probability distribution of these variables yields the usage model.

Together, device and usage models together describe a compact model for residential electrical loads. Below, we discuss the key steps in device and usage modeling in detail.

4.2 Device Modeling

Figure 5 depicts the key steps for automated device modeling, which are outlined in the previous section. We discuss each step in more detail below.

4.2.1 Active Period Extraction

As noted earlier, each load alternates between active and inactive periods. During its inactive period, where the load is off or in standby mode, the load will either consume zero power or a small amount of standby power (also known as "vampire" power [11]). Hence, given a trace of raw power usage, inactive periods can be determined by sequentially scanning the trace for periods where the power usage is less than a low threshold ϵ for durations longer than a threshold interval τ. Once inactive periods are labeled in the trace, the remaining periods are, by definition, active periods. This step partitions the power trace of the load into segments of active and inactive periods.

4.2.2 State Change Detection

When a load is active, it may transition between different active states. Each state may represent transitions between different basic loads that are components of the overall load, or may represent different active states of a basic load. Each state manifests itself in terms of a different power usage pattern. For example, a washing machine cycle may involve wash, rinse, and spin cycles, where different components of the washer (i.e., basic loads) activate in turn. Similarly, during the spin cycle, the motor may transition through different speeds, each of which is a distinct state with a different power usage level.

Since each active state has a distinct and observable power usage pattern, our technique uses a *change point detection algorithm* to determine when significant changes (i.e., transitions) occur in the observed power usage. Change point detection (also known as change detection) is a well-known technique that is used for anomaly detection [18, 21]. However, since traditional change detection techniques are not well suited to our problem, we devise a new change detection algorithm to detect state transition points within an active period.

Our energy-specific change point detection algorithm is based on the notion of *approximate entropy*. Intuitively, entropy is a measure of the unpredictability of information content. In the context of time series data, *Approximate Entropy* (ApEn) is a technique to quantify unpredictability of fluctuations in data [19]. Our algorithm operates over a sliding window of the power time series for an active period. For each position of the sliding window, it computes the approximate entropy H over a the window of length ϕ. Next, we need to detect *large* changes in approximate entropy as the window slides over the time series. To do so, we employ the *Canny Edge Detection* algorithm [7], a technique from computer vision, to detect "edges" where there are sudden changes in the entropy values H. Further, we remove certain edges that are within a predefined range δ of each other. Doing so yields instants in the power trace where significant changes in approximate entropy (which represent active state changes) are observed. Algorithm 1 describes the pseudocode of our change detection algorithm[2] and Figure 6 illustrates the different steps in the algorithm: (i) approximate entropy computed over a sliding window, (ii) canny edge detection, and (iii) removing nearby edges for a washing machine power trace.

Given the change points, our technique then partitions each ac-

[2]The ApEn computation requires us to set two additional parameters (sequence length, set to $M = \phi/4$ and filtering level, set to $R = .2 \cdot \sigma(X)$) that are not shown in the pseudocode.

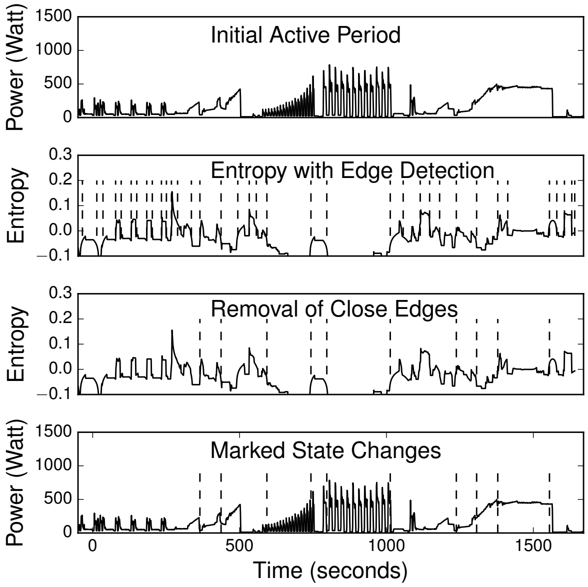

Figure 6: Approximate Entropy based change detection using canny edge detection on a washing machine trace

tive period into segments, where each segment represents the power usage observed in a specific active state.

Algorithm 1 Changepoint detection to mark active state changes using Approximate Entropy

procedure ENTROPY-CHANGEPOINT(X, ϕ, δ)
Initialize: $H \leftarrow []$
$H.append(ApEn(X[i : i + \phi])) \quad \forall i \in [1..|X|] - \phi]$
$\varepsilon_{all} \leftarrow CannyEdge1D(H)$
$\varepsilon \leftarrow RemoveCloseEdges(\varepsilon_{all}, \delta)$
return ε

4.2.3 Cycle Detection

Certain loads may transition through repeating cycles of active states, yielding cyclic behavior that manifests itself as repeating patterns of observed power usage. Hence, rather than modeling the load as a linear sequence of active states, we search for repeating sub-sequences of active states that represent cyclic behavior within each active period or repeating patterns within an active state. We use *autocorrelation*, a standard time series technique, to discover repeating power patterns within an active period. The autocorrelation of a periodic signal will exhibit a *local maxima* at the time multiples of the original signal's underlying period. Thus, we compute the autocorrelation of the active period time series for different lag values to determine cycles. To illustrate this process, we choose a portion of the washing machine trace and show the corresponding autocorrelation values for the identified cycles (see Figure 7).

4.2.4 Model Fitting

After extracting a time series segment for each active state within an active period, our technique then turns to the key problem of deriving an analytic model that describes the power usage variations observed within each state. Recall from Section 3 that a basic load can exhibit on-off, on-off decay, on-off growth, stable-min or stable-max behavior, depending on whether it is resistive, inductive or non-linear. We use analytic closed form equations to capture the behavior of the first three types of loads and use probability distributions to capture the behavior exhibited by latter two types of non-linear loads. The model fitting process involves fitting a curve onto the time series data for the first three load types and fitting a distribution onto the data for the latter two. Since we have no a priori knowledge of the load type, our approach tries to fit different types of curves or distribution and chooses the best fit.

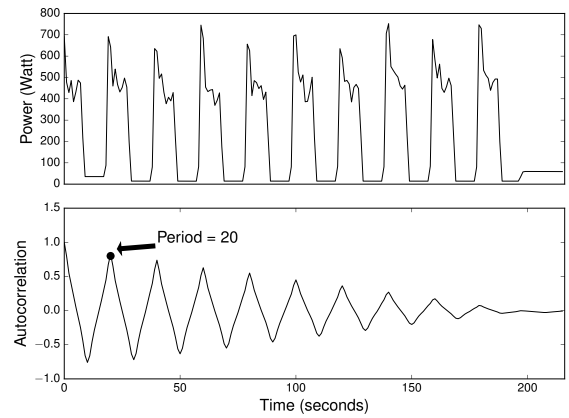

Figure 7: Autocorrelation plot of a time segement in an active duration of a device

First, to determine whether to fit a curve or a distribution, our technique determines if there are noticeable trends in the data (i.e., on-off, on-off decay or growth) or if it is derived from a random process (stable min or stable max). This is achieved by differentiating the time series of the load and observing the change in standard deviation. Our insight is that the standard deviation of the differentiated time series should decrease for trending data and increase for data derived from a random process. This step enables our technique to determine whether to fit a curve or fit a distribution for each time series segment corresponding to an active state.

In the former case, our technique then attempts to fit a linear segment, an exponential decay curve and logarithmic growth curve onto the data using *non-linear least squares* method. In the latter case, our technique attempts to fit both the gamma and the log-gamma distributions onto the data using the *Maximum Likelihood Estimation* (MLE) method. In either case, we choose the curve or the distribution that is the best fit in terms of explaining the observed data. Specifically we use *goodness-of-fit* measures, discussed in Section 6, to choose the best fit. The output of this step is a classification of each active state as a particular type of base load and the parameters of the derived model (i.e., curve or distribution) for that base load. Figure 8 illustrates the *on-off decay* fit on the part of the time segment shown in Figure 7.

4.2.5 Device Model Generation

The previous step derives a unique model for each non-repeating active state within an active period and repeats this process for each active period in the raw time series. This yields a collection of models and our final step derives an overall device model from this collection of base models. This is achieved by creating a multi-tuple record comprising models for each active period. Each tuple contains information on the state number (in a given active period), period (or 0 if no period is found), the chosen label for the model (on-off decay, stable max etc.), the fit parameters (P_{stable}, P_{peak}, λ and time length for on-off decay), and overall segment length. For the segment shown in Figure 7, for instance, the tuple ⟨Segment

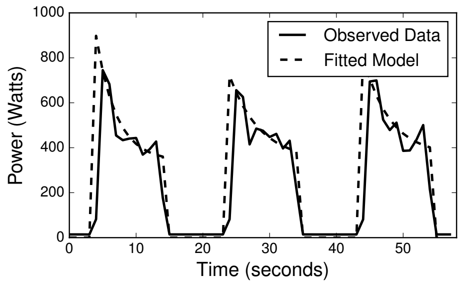

Figure 8: *On-Off Decay* model fit on a time segement in a active duration of a device

Peak no.	P_{stable}	P_{peak}	λ	timelength
1	339.78	897.85	0.33	12
2	342.68	719.03	0.22	12
3	366.87	805.28	0.25	12
Mean	349.77	807.38	0.27	12

Table 1: Variation in parameters of the active duration of a device shown as a frequency table with mean value.

Number, Period, Model, Fit Parameters, Segment length ⟩ is given by - $\langle 11, 20, On-off \quad Decay, *params, 200\rangle$.

In the case of cycles within an active period or when the same active state is observed *across* active periods, the same basic model will be found repeatedly. However, due to the power behavior of electrical loads, there may be slight differences in the observed power values or patterns for different observed instances of the same state. Hence, the computed parameters of the load will vary slightly from one instance of the state to another. Our overall model can capture the variations at different degrees of accuracy. A more accurate description is less compact but captures the observed variations more faithfully. Conversely, a more compact model is less accurate and also more approximate. Currently, our technique supports three representations for capturing parameter variations across repeating instances of the same active state: (i) a single mean value for each parameter across all instances (the most concise representation, but also the least precise), (ii) a frequency table, or (iii) a probability density as multiple dimensions in a parameter hyperspace (the most precise). Table 1 displays the frequency table and the mean value for all the parameters for the time segment shown in Figure 7. Note that a single mean value for each model parameter will lose subtle variations exhibited by the load, while a probability density captures the likelihood of all the possible values of the different parameters of a model.

4.3 Usage Modeling

The usage model captures how a load is used within a certain environment by its users. Regardless of whether the load is a foreground load or a background load, the usage of a load is captured by how frequently it activates and when. Hence, the usage patterns can be captured by three parameters: (i) start time, (ii) length of an active period, and (iii) length of an inactive period. Note that the three parameters are not independent—the end of an inactive period defines the start time of the next active period. Nevertheless deriving all three parameters enables us to capture both the frequency of usage as well as seasonal dependencies (e.g., load is only active in the evenings, or only on weekends, or only in the summer etc). Figure 9 shows energy consumption (in kWh) in the form of a heat map for an AC, a clothes dryer and a refrigerator for each hour of the day for an entire year. The figure shows that the AC is used predominantly in the summer, while the refrigerator is active multiple times every single day on account of being an "always-on" load. The dryer is typically used only once or twice a week.

Step 1: To capture various usage patterns, our technique first determines the smallest time window (e.g., day, week, month or year) over which the load exhibits statistically significant usage variations. To do so, we start with the largest time window present in the trace (e.g., a year or a month) and compute the frequency distribution of start times over this time window. We then compute the *coefficient of variation* for the start time frequency ν, given by

$$c_v = \frac{\sigma(\nu)}{\mu(\nu)} \tag{5}$$

We then recursively proceed to the next smaller time window (e.g., pick a week if the previous window was a month) and repeat the process of computing the frequency distribution of start times over this window and the coefficient of variation (COV) until the COV is found to be greater than 1. In doing so, we pick the smallest time window (i.e., the most compact temporal representation) to model usage while ensuring that we do not miss any statistically significant variations in usage of the load. In the example shown in Figure 9, this would yield periods of a day and a week, respectively, for the refrigerator and the dryer. Although, seasonal changes are captured using this, we do not directly incorporate temperature as a parameter in our model generation.

Step 2: Given the appropriate time window over which usage should be modeled, our technique then uses the start times and lengths of active and inactive periods extracted from the power series trace to compute (i) a histogram of start times over the time window, and (ii) histograms of active and inactive period lengths. We then use the Kernel Density Estimation (KDE) approach to compute a smooth probability distribution over each histogram. KDE is a non-parametric method for data smoothing when one needs to reason about the population based on limited samples. This process yields three probability distribution functions for the start times, active and inactive period lengths, respectively. The joint probability distribution function over these three parameters represents the usage model for the load.

5. NIMD IMPLEMENTATION

We implemented a prototype of our NIMD algorithm in python using the SciPy stack. SciPy stack has a collection of powerful scientific computing libraries for data processing. Our prototype takes a raw power trace as input and outputs a device model and the usage model for it using techniques described in the previous section. The overall model fitting component and Kernel Density Estimation uses specific modules from the SciPy library. For calculating Approximate Entropy, we used PyEEG [2], an open source python module for data processing for EEG data. For other statistical mechanisms, we used standard python libraries.

The model derived from the trace can then be employed for a number of higher level energy algorithms. In addition, the model, which is a compact description of the device, can be also used to create a synthetic traces that “faithfully” mimic the load’s actual power behavior, as discussed next.

6. EVALUATION

In this section, we evaluate the efficacy of our NIMD approach

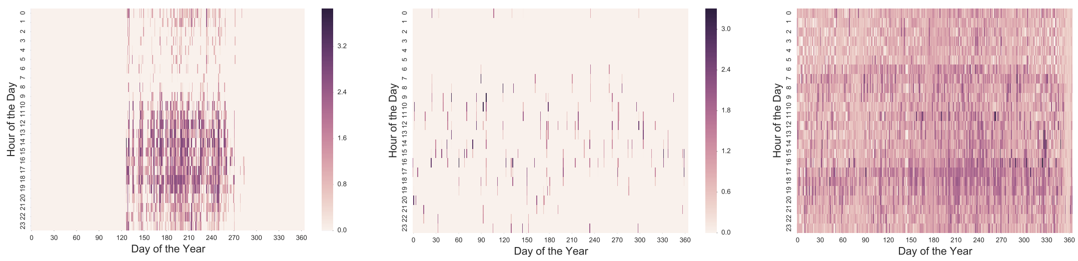

(a) Airconditioner (b) Dryer (c) Refrigerator

Figure 9: Energy consumption of devices in kWh over different time of the day and day of the year

Name	#Devices	Duration	Frequency	Region
AMPds	24	2 years	1 Minute	Canada
Smart*	26	3 months	1 second	USA
Tracebase	158	few days	1 second	Germany

Table 2: Datasets used for evaluation

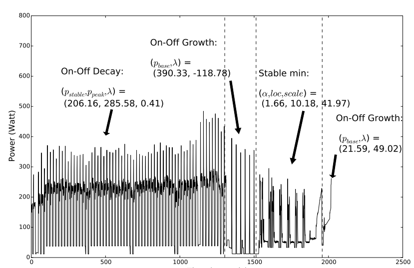

Figure 11: Model learnt for a composite load

for device and usage modeling.

Datasets: We used device-level electrical data from three publicly available datasets: AMPds [17], Smart* [5], and Tracebase [20]. Table 2 describes the key characteristics of these datasets. AMDds is the smallest of the three datasets, but has load data for two years. Tracebase is the most extensive dataset in terms of number of loads, while the Smart* has appliance-level data at a 1-second resolution over a period of 3 months.

Metrics: To analytically evaluate the goodness of fit for *on-off*, *on-off decay* or *on-off growth* models described earlier, we use Mean Absolute Percentage Error (MAPE), a standard statistical measure of accuracy expressed as a percentage value. The formula for calculating MAPE for a given device with power consumption data represented as $X^{data}_{[1..k]}$ and the fitted model $X^{fit}_{[1..k]}$ is given below.

$$MAPE = \frac{100}{n} \cdot \sum_{k=1}^{K} \left| \frac{X^{data}_k - X^{fit}_k}{Mean(X^{data}_{[1..K]})} \right| \qquad (6)$$

For *stable min* and *stable max* models, we use Kullback-Leibler (KL) divergence, a measure of the difference between two probability distributions. KL divergence of probability distribution Q from P is symbolized as $D_{KL}(P||Q)$. However, KL divergence is not a metric as it is not symmetric. In practice, probability P is the distribution of the data and Q is the proposed approximation for P. In our case, Q is the $Gamma$ distribution for *stable min* and the $LogGamma$ distribution for *stable max*. The lower the MAPE or KL divergence values, better is the fit.

$$D_{KL}(P||Q) = \sum_i P(i) \cdot \log \frac{P(i)}{Q(i)} \qquad (7)$$

6.1 Device Modeling

6.1.1 Model Fitting

Figure 10 illustrates the performance of model fit on the 4 appliances from the Smart* dataset. These appliances are - (a) a Refrigerator, (b) an AC, (c) a CRT-Monitor, and (d) a LCD-TV fitted with *on-off decay*, *on-off growth*, *stable min*, and *stable max* models respectively. The learnt model parameters are also shown in each figure. Figure 10 (a) and (b) show the MAPE values for the fitted model on the data. For the two examples shown for curve fits in Figure 10(a) and (b), we get a MAPE (error) of 2.4% and 1.02%. Figure 10 (c) and (d) show the KL divergence of data from $Gamma$ and $LogGamma$ distribution respectively. These figures also show the KL divergence of data for a baseline $Normal$ distribution fit. Intuitively, KL divergence is the penalty on compressing data to be represented as the proposed distribution. The figure shows that KL divergence of the our proposed $Gamma$ and $LogGamma$ distributions is a more than a factor of 2 lower than the baseline $Normal$ distribution. Finally, Figure 11 shows the overall model learnt for a washing machine, a composite load.

6.1.2 Accuracy of the models

To evaluate the accuracy of model fit, we ran it on a number of appliance loads of various types in the tracebase dataset. In Figure 12(a), we have a violin plot showing the MAPE values for 5 refrigerators over curve fit on several active periods of the device. The horizontal stick in these plots represents each underlying datapoint corresponding to a measurement for an active period. The thickness of the graphs for different devices corresponding to the MAPE values on the y-axis is indicative of the frequency distribution of the datapoints. Overall, more than 1000 active periods spread across 5 devices are shown in this figure. MAPE values for 2 of the refrigerators are almost below 3%, whereas it is between 1-7% for 2 other refrigerators. For one refrigerator, we found a comparitively much poorer fit in the range of 6-10%. Figure 12(b) shows an appliance type-wise view of the error in the curve fits for 4 inductive (Refrigerator) or resistive (Kettle, Lamp, Toaster) load types. The graph represents more than 1300 active period data for 7 Lamps, 6 Kettles and 2 Toasters along with the 5 refrigerators shown in (a). As shown, the resistive loads have MAPE values below 4%.

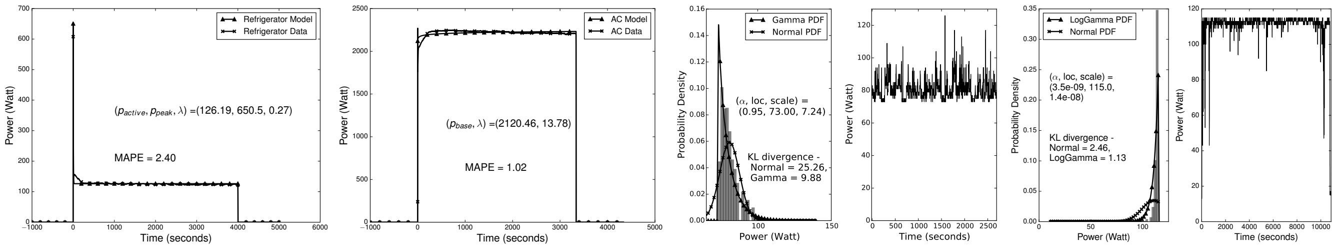

(a) Refrigerator (b) AC (c) CRT-Monitor (d) LCD-TV

Figure 10: Basic Load Models for *On-off Decay*, *On-off Growth*, *Stable min*, and *Stable max* with fitted models

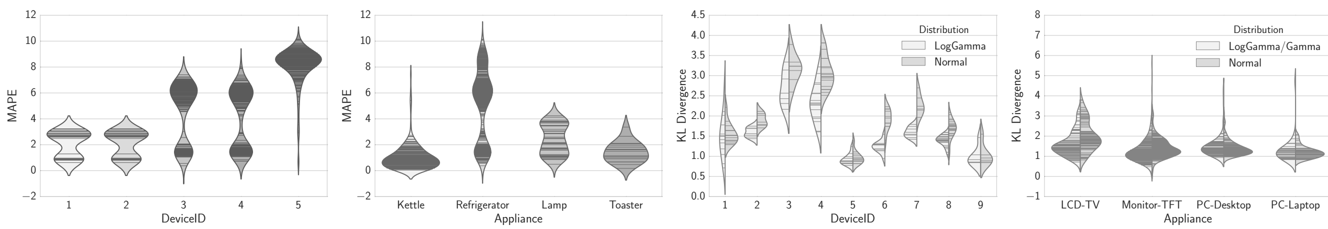

(a) Refrigerators (b) Resistive and Inductive Loads (c) LCD-TVs (d) Non-linear Loads

Figure 12: Goodness of Fit measures for different appliances from Tracebase dataset

As discussed earlier, for distribution fit we use a relative measure called KL divergence. Here, again we compare our proposed distributions to a baseline $Normal$ distribution. Figure 12(c) shows the violin plots for more than 200 active periods spread across 8 LCD-TVs. For 5 devices the KL divergence improves by modeling the active period traces as a $LogGamma$ distribution by a factor of 1.5. For the other 3 devices, there is no appreciable difference between the two distributions. Figure 12(d) represents appliance type-wise spread of KL divergence for non-linear loads such as TFT-Monitor, Desktop-PC, and Laptop along with LCD-TVs representing more than 1000 active periods. Except for LCD-TVs, we do not find any improvement (or worsening) in KL divergence by model fitting proposed one-tailed distributions over $Normal$ distribution.

6.1.3 Descriptiveness of the Model

The model parameters of an electrical load are not static and the variation in them must be captured for building a realistic model. Figure 13 shows the probability density over the 3 dimensions of the parameter space (p_{peak}, p_{stable}, and λ) obtained from applying NIMD algorithm on the different active periods of a refrigerator from the TraceBase dataset. We observe that the 3 parameters vary from one active period to the other. Figure 13 illustrates how a probability density is more precise than a frequency distribution table (shown as a scatter plot) as it provides a smooth parameter space with just a few data samples.

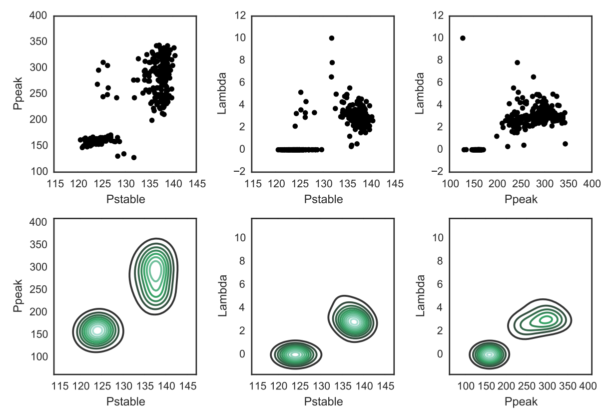

Figure 13: Variation in parameters over several active periods

6.2 Usage Modeling

To evaluate the efficacy of the usage modeling, we used loads from the AMPds dataset since it contains consumption data for a period of 2 years. With an adequate amount of data, we can choose the smallest time window which captures the usage variations of a device. In the section 4, we discussed how the joint probability distribution of start times over an optimal window and the length of the active and inactive periods capture the usage model of any device. Since the joint probability of these 3 variables is difficult to plot, we use Table 3 to show the mean and the standard deviation of a number of active periods in a time window (optimally selected from the data) with the length of the active and inactive periods. The time window selected for the different devices matches the intuitive values that would have been manually selected for the different devices. For example, our models indicate that the approximate duty cycle for the refrigerator is around 36 minutes (average active + inactive period lengths). Our usage models also capture that devices such as Dryers, Washing machines and Dishwashers are used 3 to 6 times per week.

6.3 Automated versus Manual Modeling

To compare our automated approach with models manually derived by experts, we obtained load data and manually derived models reported in [4] from the authors. We used NIDM to derive models for the loads and then compare NIMD's models with the manually derived ones. Figure 14 shows a comparison between the manual and the automated modeling approaches. We were able to classify each of these 4 appliances with the correct basic load type. Further, the learnt parameters were very close to the manual modeling values shown in [4]. The error associated with both manual and automated modeling was lesser than 1% in all 4 cases. Thus, our automated approach derives models comparable to human-derived

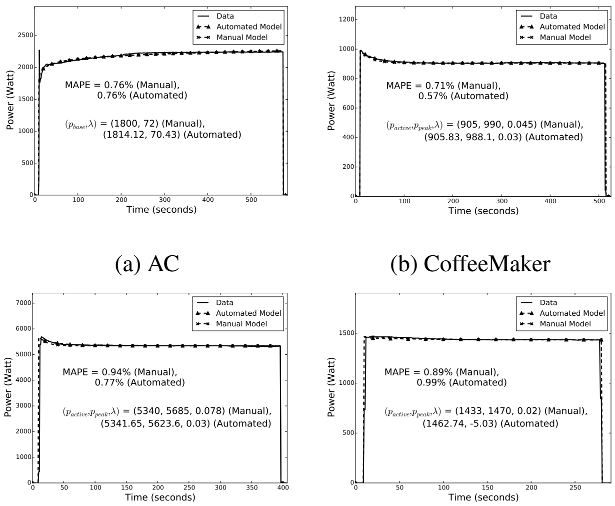

(a) AC (b) CoffeeMaker

(c) Dryer (d) Toaster

Figure 14: Comparison of automated v/s manual modeling

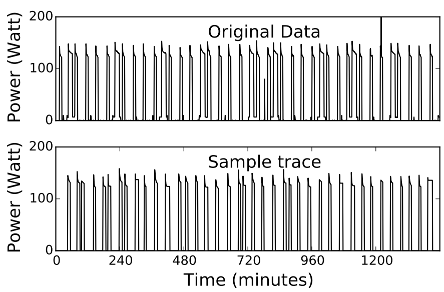

Figure 15: Synthetic trace generated using out models.

fitted models using domain knowledge (e.g., load type).

6.4 Case study: Synthetic Trace Generation

While our models can be used for many energy management tasks, they can also be used to derive a synthetic power trace that is statistically similar to the original trace. For this, we need to sample the usage distributions to compute start times of each active and inactive durations. To derive the parameters such that the synthetic trace mimic the original load usage, we need to draw samples from joint probability distribution computed by the usage model. To do so, we employ a state-of-the-art Markov Chain Monte Carlo sampling method called the *Metropolis-Hastings Algorithm* [10] that generates a sequence of samples through a *random walk* over the sample space. Once the start time is computed, the usage model, which is itself a sequence of analytic models for each active state, is then used to create a power trace for that active period. In the case of non-linear loads, where the device models are distributions, we sample the distribution to create a power trace. At the end of an active period, we set the power to the standby power level for the inactive period. The process repeats for the next start period.

Figure 15 shows the original load trace and the sample trace generated synthetically from the first sample taken from real data after 10^4 iterations. By "replaying" the usage and device models, we observe that the synthetic trace exhibits similar power usage as the original load trace.

Device (Window)	Start times/interval		Active length		Inactive length	
	μ	σ	μ	σ	μ	σ
Dryer(W)	4.2	2.2	41.8	11.7	1891	2220
Washer(W)	5.2	2.7	50.9	24.3	1557.1	2114.7
Dishwasher(W)	3.7	1.3	75.8	39.6	2034	2160
Fridge(D)	39.6	4.0	13.5	9.8	22.5	9.7
TV(D)	1.9	1.0	67.7	47.9	522.8	501.5
WOE(M)	3.6	1.8	105.9	399.6	8978	7901

Table 3: Usage Patterns for different devices with Daily (D), Weekly (W) or Monthly (M) window

7. RELATED WORK

Due to the large-scale deployment of smart meters by utilities, there has been a resurgence in interest in energy analytics techniques, such as NILM, in both academia [12, 1, 16] and industry [8]. NILM-based energy analytics have been used in different scenarios, such as opportunistic load scheduling for capping peak demand [6], learning thermostats schedule [13], etc. However, prior work on NILM generally uses simple *on-off* models for electrical loads, which, as we show, are highly inaccurate. As a result, these techniques have limitations on their accuracy.

Thus, an important challenge is the ability to analytically model the behavior of a variety of residential loads. Earlier work [3, 4] has demonstrated that most appliances map onto few basic types that exhibit a compact set of features. This prior work shows how to manually construct models for the basic load types, but does not show how to automatically derive models, especially for complex loads that are time-consuming to manually model.

In this work, we propose an algorithm to automatically derive a model for each appliance from its empirical measurements. Our technique is analogous to disaggregation where an energy usage trace of a compound load is automatically disaggregated into a set of basic load types and the parameters of each basic load type are automatically learned. Further, we also model the interaction of devices with residents to build a usage model for them. Finally, our NIMD techniques adapt and extend multiple methods from probability, statistics, and information theory to the energy analytics domain. These methods provide a strong theoretical framework for automatically deriving models of electrical load behavior.

8. CONCLUSIONS

In this paper, we presented a new approach for automated unsupervised derivation of the device and usage models of residential loads. We presented our NIMD approach that uses concepts from power systems, statistics, and machine learning to automate loads modeling. Our experimental evaluation showed that our automated models are within 1% of the ground truth and very close to those derived manually by experts and yield good fits for a range of loads. A current limitation of our approach is that they only handle sequential composite loads, where the base loads activate in sequence, and do not handle parallel composite loads. As future work, we will study methods that combine NILM disaggregation with our NIMD approach to handling parallel composite loads.

Acknowledgements. This research is supported by NSF grants IIP-1534080, CNS-1405826, CNS-1253063, CNS-1505422, and the Massachusetts Department of Energy Resources.

9. REFERENCES

[1] K. Armel, A. Gupta, G. Shrimali, and A. Albert. Is Disaggregation the Holy Grail of Energy Efficiency? The Case of Electricity. *Energy Policy*.

[2] F. Bao, X. Liu, and C. Zhang. PyEEG: An Open Source Python Module for EEG/MEG Feature Extraction. *Computational Intelligence and Neuroscience*, 2011.

[3] S. Barker, S. Kalra, D. Irwin, and P. Shenoy. Empirical Characterization and Modeling of Electrical Loads in Smart Homes. In *IGCC*, July 2013.

[4] S. Barker, S. Kalra, D. Irwin, and P. Shenoy. Empirical Characterization, Modeling, and Analysis of Smart Meter Data. *IEEE Journal on Selected Areas in Communications*, 32(7), July 2014.

[5] S. Barker, A. Mishra, D. Irwin, E. Cecchet, P. Shenoy, and J. Albrecht. Smart*: An Open Data Set and Tools for Enabling Research in Sustainable Homes. In *SustKDD*, 2012.

[6] S. Barker, A. Mishra, D. Irwin, P. Shenoy, and J. Albrecht. Smartcap: Flattening Peak Electricity Demand in Smart Homes. In *PerCom*, March 2012.

[7] J. Canny. A computational approach to edge detection. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, PAMI-8(6), Nov 1986.

[8] M. Capps. Energy and IOT. An Engineer's Perspective. *XRDS*, 22(2), December 2015.

[9] D. Chen, S. Barker, A. Subbaswamy, D. Irwin, and P. Shenoy. Non-Intrusive Occupancy Monitoring using Smart Meters. In *BuildSys*, November 2013.

[10] S. Chib and E. Greenberg. Understanding the Metropolis-Hastings Algorithm. *The American Statistician*, 49(4), November 1995.

[11] M. Ebling and M. Corner. It's all about power and those pesky power vampires. *IEEE Pervasive Computing*, 8(1), 2009.

[12] G. Hart. Nonintrusive Appliance Load Monitoring. *Proceedings of the IEEE*, 80(12), December 19 92.

[13] S. Iyengar, S. Kalra, A. Ghosh, D. Irwin, P. Shenoy, and B. Marlin. iProgram: Inferring Smart Schedules for Dumb Thermostats. In *BuildSys*, November 2015.

[14] J. Kelso, editor. *2011 Buildings Energy Data Book*. Department of Energy, March 2012.

[15] W. Kleiminger, C. Beckel, T. Staake, and S. Santini. Occupancy Detection from Electricity Consumption Data. In *BuildSys*, November 2013.

[16] J. Kolter and A. Ng. Energy Disaggregation via Discriminative Sparse Coding. In *NIPS*, December 2010.

[17] S. Makonin, F. Popowich, L. Bartram, B. Gill, and I. Bajic. AMPds: A Public Dataset for Load Disaggregation and Eco-Feedback Research. In *EPEC*, October 2013.

[18] D. Picard. Testing and estimating change-points in time series. *Advances in Applied Probability*, 17(4), 1985.

[19] S. Pincus. Approximate entropy as a measure of system complexity. *Proceedings of the National Academy of Sciences*, 88(6), 1991.

[20] A. Reinhardt, P. Baumann, D. Burgstahler, M. Hollick, H. Chonov, M. Werner, and R. Steinmetz. On the Accuracy of Appliance Identification Based on Distributed Load Metering Data. In *SustainIT*, October 2012.

[21] X. Shao and X. Zhang. Testing for change points in time series. *Journal of the American Statistical Association*, 105(491), 2010.

[22] Belkin wemo. http://www.wemo.com/, February 2016.