PEAT, how much am I burning?

Nambi, S. N.Akshay Uttama; Prasad, R. Venkatesha; Lua, Antonio Reyes; Gonzalez, Luis

DOI
10.1145/3204949.3204951

Publication date
2018

Document Version
Final published version

Published in
Proceedings of the 9th ACM Multimedia Systems Conference, MMSys 2018

PEAT, *How Much Am I Burning?*

Akshay Uttama Nambi S. N., Venkatesha Prasad R, Antonio Reyes Lua and Luis Gonzalez
Delft University of Technology (TUDelft), The Netherlands
akshay.uttama@gmail.com,R.R.VenkateshaPrasad@tudelft.nl

ABSTRACT

Depletion of fossil fuel and the ever-increasing need for energy in residential and commercial buildings have triggered in-depth research on many energy saving and energy monitoring mechanisms. Currently, users are only aware of their overall energy consumption and its cost in a shared space. Due to the lack of information on individual energy consumption, users are not being able to fine-tune their energy usage. Further, even-splitting of energy cost in shared spaces does not help in creating awareness. With the advent of the Internet of Things (IoT) and wearable devices, apportioning of the total energy consumption of a household to individual occupants can be achieved to create awareness and consequently promoting sustainable energy usage. However, providing personalized energy consumption information in real-time is a challenging task due to the need for collection of fine-grained information at various levels. Particularly, identifying the user(s) utilizing an appliance in a shared space is a hard problem. The reason being, there are no comprehensive means of collecting accurate personalized energy consumption information. In this paper we present the *Personalized Energy Apportioning Toolkit (PEAT)* to accurately apportion total energy consumption to individual occupants in shared spaces. Apart from performing energy disaggregation, PEAT combines data from IoT devices such as smartphones and smartwatches of occupants to obtain fine-grained information, such as their location and activities. PEAT estimates energy footprint of individuals by modeling the association between the appliances and occupants in the household. We propose several accuracy metrics to study the performance of our toolkit. PEAT was exhaustively evaluated and validated in two multi-occupant households. PEAT achieves 90% energy apportioning accuracy using only the location information of the occupants. Furthermore, the energy apportioning accuracy is around 95% when both location and activity information is available.

CCS CONCEPTS

• Information systems → Mobile information processing systems; **• Human-centered computing → Ubiquitous and mobile computing design and evaluation methods**; **• Hardware → Energy distribution**; **Smart grid**;

MMSys'18, June 12–15, 2018, Amsterdam, Netherlands

ACM ISBN 978-1-4503-5192-8/18/06…$15.00
https://doi.org/10.1145/3204949.3204951

KEYWORDS

Energy apportioning, NILM, shared spaces, IoT

ACM Reference Format:
Akshay Uttama Nambi S. N., Venkatesha Prasad R, Antonio Reyes Lua and Luis Gonzalez. 2018. PEAT, *How Much Am I Burning?*. In *MMSys'18: 9th ACM Multimedia Systems Conference, June 12–15, 2018, Amsterdam, Netherlands.* ACM, New York, NY, USA, 12 pages. https://doi.org/10.1145/3204949.3204951

1 INTRODUCTION

Global energy consumption in residential and commercial buildings is estimated to be 35-40% of generation. This is expected to rise further due to increase in a number of appliances and electronic devices [1]. Many investigations, hitherto, focused on real-time energy monitoring in households and also provided feedback to the occupants about their energy consumption [2–5]. These studies have shown that feedback to occupants on their energy consumption in real-time can effectively raise awareness and promote energy conservation behavior [8]. With the advent of the Internet of Things (IoT), multiple data streams from various sensors could be gathered. Thus the corroborating information about the location of a user and the appliance being used can provide much more dividend. This work explores a focused and highly difficult challenge that is providing individual energy footprint with least number of sensors in shared spaces.

To provide users with detailed energy consumption breakdown several intrusive [3–6] and non-intrusive methods [9, 10] have been proposed. However, Non-Intrusive Load Monitoring (NILM) techniques have prevailed due to its lower deployment cost. NILM aims to estimate appliance level energy consumption from the aggregate consumption data of households. While these techniques help in understanding energy consumption in a building, they lack the ability to provide energy footprint of individuals. A recent study highlighted advantages of providing energy consumption information to individuals and has the potential to reduce up to 20% of the total energy consumption [11, 12]. In shared spaces – such as student housing, office environments, multi-occupant household – lack of individual occupant (from now on, we use the terms 'individual occupant', 'individual' and 'per-occupant' interchangeably) consumption data necessitates even-splitting of energy cost. This results in inefficient energy usage where occupants minimize their own cost by taking advantage of others [11, 13]. To analyze this we collected energy consumption data in a multi-occupant student housing and a residential household. Fig. 1 shows the daily energy usage of occupants across various appliances. It can be clearly seen that the amount of energy consumed by each occupant significantly varies in both the settings. This necessitates the need for *personalized energy disaggregation* system moving towards user level from appliance level. The disaggregated individual energy usage information can be used to develop better energy management techniques apart from raising awareness [17].

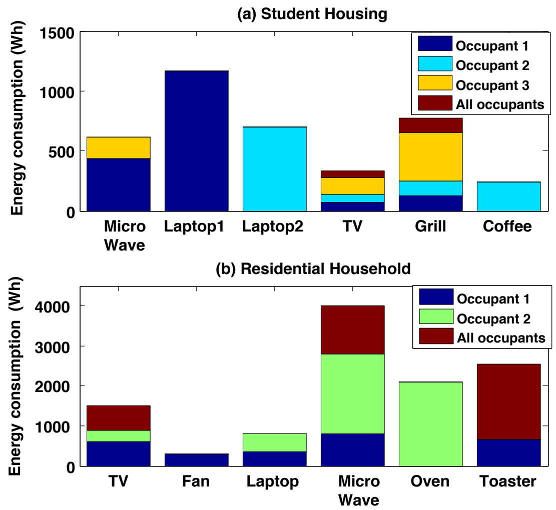

Figure 1: Energy consumption distribution across occupants.

The current literature focuses mainly on energy consumption of buildings or households as a whole. NILM techniques fail to provide energy disaggregation at the user level, limiting the opportunities for encouraging energy-saving behavior. Aggregated energy information of households fail to answer questions such as,*"how much energy each occupant has used today?", "who amongst us burns most of the energy?", "which occupant is more energy efficient in the household?"*. Currently, there are no comprehensive means of providing information to individual occupants on their energy consumption. Hence, it is required to an design energy apportioning system that disaggregates the total consumption of a household into per-occupant, per-appliance level.

Recently, there have been few research efforts [14, 15, 18–20] that aim to study per-occupant energy footprint in various settings. However, providing energy consumption information of individuals, in shared spaces, is a challenging task due to the following requirements: (i) Collection of fine-grained information at various levels. This fine-grained information include *which* occupant performed *what* activity and *where* in the household. (ii) Deployment of additional sensors to identify the occupant using a particular appliance at a specific location. This is cumbersome in terms of cost and maintenance. (iii) Real-time identification of the correct occupant using an appliance when multiple occupants are present in the same location. The current approaches apportion the energy equally to all the occupants in that location [14, 19]. (iv) Most of the techniques, proposed hitherto, are centralized with either third-party services or energy utilities having privacy-sensitive information of consumers. This raises several issues related to scalability and privacy. (v) The resolution of apportioning may vary depending upon the environment i.e., residential household -Ű where some of the shared appliances is for total family consumption (e.g., kitchen utilities, fridge); student housing Ű - where a shared appliance is used by only one occupant at that instance (This can also be seen in Fig. 1). Hence, identifying the granularity of apportioning is a challenging issue.

To this end, we present **P**ersonalized **E**nergy **A**pportioning **T**oolkit (PEAT) that combines readily available data from the ubiquitous IoT sensors present in the household to derive fine-grained occupant level energy consumption information. Specifically, we use a single smart meter energy data to derive fine-grained appliance level energy information. A modified Combinatorial Optimization (ModCO) algorithm is proposed to significantly reduce the computational complexity for energy disaggregation. Internet of things (IoT) devices such as smartphones and/or smartwatches with WiFi radios are used for indoor room level localization and to determine the activities performed by the occupants. In this work, data from these sensors are used to detect user occupancy and the *micro-activities* performed by occupants. PEAT combines NILM technique with WiFi-based localization and activity monitoring to determine *when* an appliance is being used, and *which* occupant is currently using the appliance. Specifically, PEAT studies and models the association between appliances and occupants in a shared space. Furthermore, PEAT runs on a low-cost embedded system such as Raspberry Pi, to apportion energy consumption to individuals. Privacy-sensitive data of occupants are stored and processed locally at the household making this approach highly scalable and privacy preserving. Our system was extensively evaluated in two real-world multi-occupant settings, *viz.*, (i) student housing and (ii) residential house. The collected data and the toolkit are publicly available[1] for the community to support additional analysis. The main contributions of this paper are:

- We present a personalized energy apportioning toolkit (PEAT) to derive real-time per-occupant energy footprint in shared spaces.
- We propose a modified CO algorithm to accurately derive fine-grained appliance level information from aggregated energy consumption.
- We describe our inference algorithm, which models the association between appliances and users to study the dynamics of appliance usage.
- We provide an extensive experimental evaluation of PEAT from two multi-occupant buildings, *viz.,* student housing and residential settings. PEAT was empirically evaluated in both the settings for over two months.

2 RELATED WORK

Energy apportioning in shared spaces covers a broad range of research areas from appliance monitoring, user monitoring to personalized energy monitoring. Recent studies have shown that providing per-occupant energy footprint offers the potential for energy reduction and promotes energy saving behavior [2, 11, 12]. We first describe in detail the state-of-the-art techniques proposed for monitoring in shared spaces.

Appliance monitoring. Energy consumption of a household at various levels can be derived either by embedding more sensors in the household [3, 4, 6, 7] or by using a minimal number of sensors [9, 10]. Ho et al. [7] employ a thermal camera to detect a change in appliance surface temperature to infer appliance on/off state. However, their system can only detect on/off state of

[1]DRED dataset: http://www.st.ewi.tudelft.nl/~akshay/dred/

Table 1: Comparison of various efforts towards energy apportioning.

Work	Environment	Architecture	Appliance monitoring	User monitoring	Apportionment policies	Additional sensor
[14]	Office	Centralized	Appliance level	Security access logs	Cannot distinguish between occupants	NA
[15]	Residential	Centralized	Appliance level	room level localization using doorway sensors	Evenly apportions to all occupants present in the room for shared appliances	Doorway sensors
[18]	Office	Centralized	Appliance level	proximity based using magnetic inductance	Assigns usage to the nearest occupant	Temperature, special wearables
[19]	Residential	Centralized	Audio based appliance detection	WiFi based room level localization	Evenly apportions to all occupants present in the room for shared appliances	Microphone for appliance detection
[20]	Residential	Centralized	Appliance level with unique sensors	RFID anklets and RFID antennas	Applies heuristics to determine per-occupant footprint	RFID antennas
PEAT	Office & residential	Distributed	Aggregate data using NILM	WiFi based localization, smartwatch for activity recognition	Apportions energy based on historic consumption pattern & activities performed	None

an appliance and hence apportioning may not be accurate. Specifically, NILM techniques propose disaggregating appliance level information from a single energy meter data. *Unsupervised NILM* techniques use no prior knowledge of the appliances but often require appliances to be manually labeled [9, 10]. These approaches are computationally intensive and exact inference is intractable [21]. *Supervised NILM* techniques assume that ground truth appliance level data is available to train and develop appliance models prior to performing disaggregation [3, 16]. Hart's Combinatorial Optimization (CO) algorithm identifies step changes in the aggregate electricity consumption and matches them with appliance signature database to learn the states of the appliance [16]. These algorithms require extensive training on appliance level data to model the states accurately. To reduce the computation complexity, recent energy disaggregation algorithms utilize additional data such as occupancy information to disaggregate energy in the household [22].

In this work, we propose ModCO algorithm that extends the simple CO algorithm to eliminate appliance modeling and to reduce the computational complexity. Our semi-supervised algorithm runs on a low-cost embedded system such as Raspberry Pi to provide real-time disaggregated information.

User monitoring. User occupancy detection is a crucial element in developing user-centric energy management services [23]. Several direct and indirect approaches have been proposed in the literature to derive user occupancy information [24, 25]. *Direct* approaches employ low cost sensors such as passive infrared (PIR), reed switches, RFID tags to determine occupancy information [24]. Even though these approaches are cost-effective, installing and maintaining these sensors in a household is intrusive and cumbersome. Furthermore, these techniques do not distinguish between different occupants, if the occupants have similar characteristics. *Indirect* occupancy monitoring approaches employ WiFi or Bluetooth (BT) fingerprinting using smartphones to derive room level occupancy [25]. In this work, we use the existing infrastructure like WiFi access points (AP) deployed in the household, along with smartphones to derive the current location of the occupants. We employ simple classification techniques that run locally on smartphones to model and train the data collected from the WiFi scans.

Motivated by the large scale penetration of smartwatches and their increasing sensing capabilities, we employ them to derive user activities in the household. Smartwatches provide an opportunity to identify precisely the micro-activities performed by the occupant such as opening microwave door, opening the refrigerator door, and switching on/off an appliance.

Unlike existing activity monitoring techniques that require additional sensors carried by occupants, we argue that sensors present in the smartwatch are sufficient to determine micro-activities. Moreover, with both WiFi and BT radios available on a smartwatch, one can use them for indoor localization too.

Personalized energy monitoring. Lack of per-occupant energy footprint has resulted in even splitting of energy costs and negligent energy usage in shared spaces [12]. Table. 1 provides a concise overview of state-of-the-art techniques proposed for energy apportioning. Hay et al. [14] investigate energy apportioning in an office building. They propose static and dynamic policies to apportion shared energy usage. However, these policies assign energy evenly by determining the number of people inside the building. Moreover, they used manual logs and user annotations to determine the occupancy.

In residential settings, Lee et al. [15] propose a personalized energy auditor to apportion energy with the help of smartphones and doorway sensors. They classified the appliances into "personal" and "shared". The policy for apportioning assigns the personal appliance usage to that individual, whereas the shared appliance usage was evenly split across all occupants present in the room. Moreover, the setup and installation of doorway sensors are cumbersome and intrusive. Furthermore, user characteristics (height of occupants) and appliance level metering was considered during apportioning.

Cheng et al. [18] present a model to determine the association between human activities and observed energy consumption. They use additional sensors such as LED and light sensors to determine user movements and location.

Saha et al. [19] propose mechanisms to combine smartphone data with electricity data for accurate activity detection and energy apportioning. They use WiFi based localization to determine the location of occupants and collect audio samples from the microphone continuously to determine which appliance is being used. The proposed models require extensive data collection and training for each appliance, hindering the applicability of the work in other households.

Ranjan et al. [20] map energy apportioning to a *fixture assignment* problem to determine per-occupant energy usage. Fixture assignment is done by determining the unambiguous assignments and then learning usage patterns. However, they use custom made RFID anklet and RFID antennas for indoor location tracking and deploy sensors for each appliance to find the energy usage.

In [32] authors propose a complete system to apportion the energy usage in commercial buildings. Appliance energy consumption is obtained from the building management system and user location is obtained using WiFi/BLE localization. Energy apportioning is mainly based on the location, i.e., assign the total energy consumed based on the time spent in the corresponding room. While this approach is relevant in commercial buildings where each occupant uses his/her own room, in shared spaces, multiple occupants can remain in the same room. Thus one has to determine which occupant in the room used the appliance. PEAT overcomes this by determining the micro-activities performed by the user to determine the occupant who used the appliance.

Most of the solutions proposed until now require additional sensor deployment for either appliance monitoring or occupancy detection. Moreover, they are usually centralized. PEAT extends the state-of-the-art techniques by overcoming the above-mentioned issues. We also demonstrate the effectiveness and applicability of our system in two settings, *viz.,* student housing and residential household. Furthermore, our system can be easily replicated in any other setting with minimal user intervention. To the best of our knowledge, we are the first to validate energy apportioning across multiple settings.

3 PEAT

Considering the limitations of the state-of-the-art solutions, we propose an energy apportioning toolkit that integrates smartphones, smartwatches, and smart meter data with minimal user intervention to derive real-time per-occupant energy footprint. Fig. 2 shows the system overview of PEAT. The toolkit consists of four major components: (i) *Appliance monitoring*, (ii) *User monitoring*, (iii) *Appliance-User modeling* and (iv) *Online evaluation*. We describe each component in detail below.

3.1 Appliance monitoring

Appliance monitoring component with the help of NILM algorithm determines the state of each appliance. A change in state of an appliance from "OFF" to "ON" is considered as an event trigger. Event triggers represent the appliances, which are currently being used by the occupants in the households. We first provide a brief description of the CO algorithm for energy disaggregation [16] and then propose a modified CO (ModCO) algorithm to be used in PEAT.

Combinatorial Optimization(CO): Let $\hat{y}_t^{(n)}$ be the estimated energy consumed and $y_t^{(n)}$ be the actual energy demand of each appliance n at time t. $\overline{y}_t$ represents the aggregate energy reading of

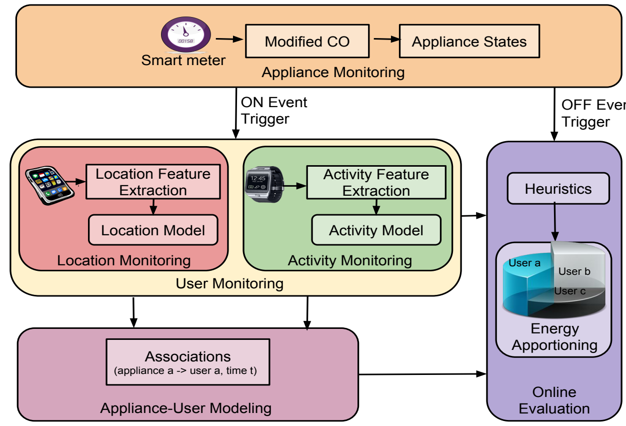

Figure 2: System overview of PEAT.

the household. The actual state of an appliance is represented by $x_t^{(n)} \in Z \geq 0$ and $\hat{x}_t^{(n)}$ represents estimated state of the appliance by the disaggregation algorithm. CO finds the optimal combination of appliance states, which minimizes the difference between the sum of power of predicted appliances and the observed aggregate power. It is given by,

$$\hat{x}_t^{(n)} = \arg\min_{\hat{x}_t^{(n)}} \left| \overline{y}_t - \sum_{n=1}^{N} \hat{y}_t^{(n)} \right| \quad (1)$$

where N is the set of all appliances in the household and t is the current time period. The predicted energy consumption of an appliance $\hat{y}_t^{(n)}$ is then mapped to the closest appliance state $x_t^{(n)}$. This approach requires models of appliances, which includes complete power consumption details for each state of the appliance. This is further used during inference step to predict the current state of the appliance.

The CO algorithm has several drawbacks. Firstly, the optimization problem resembles subset sum problem and is NP-complete. Furthermore, the computational complexity in CO increases exponentially with the number of appliances. Secondly, this algorithm does not differentiate between appliances with similar power consumption and appliances with similar states. Thirdly, this algorithm assumes all the appliances in the households are monitored and assigned some portion of energy to appliances even if they are not currently used, thus resulting in low disaggregation accuracy.

Modified CO (ModCO):

Our modified CO algorithm improves the original CO in three aspects, (i) ensures the effect of small fluctuations in aggregate power is minimal and preserves consistency in consecutive state estimations; (ii) eliminates the need of appliance level modeling by employing a crowd-sourced power consumption database, and (iii) reduces the computational complexity associated with determining the state combinations. We employ a crowd-sourced generic appliance model from the power consumption database. For example, the power consumption database[2], standby power database[3]

[2]The Power consumption database. [Online] http://www.tpcdb.com/

[3]Standby power. [Online] http://standby.lbl.gov/summary-table.html

and Oksolar database[4] provides crowd-sourced information on maximum and idle power for a wide range of loads indexed by type, manufacturer, and model number. Furthermore, our modified CO algorithm requires knowing the number of appliances and their location in the household. This metadata information is collected once during the deployment and, except a few appliances like a vacuum cleaner, the location of the appliances is generally static.

In the original CO, the algorithm tries to find the set of appliances at each time interval, which is closest to the current aggregated energy consumption. This may result in a different set of appliances being used in each consecutive interval. Hence, it is necessary to preserve consistency in selecting appliances during consecutive state estimations. We define *priority combination* – that is the set of appliances which are assumed to be currently running. This information can be retrieved from the last iteration of NILM algorithm. In each interval, ModCO first evaluates the priority combination to check whether the sum of all appliances in the priority combination matches the current aggregated value. If the difference between the sum of priority combination and the aggregated energy is within a threshold δ, then the current priority combination is retained as the predicted set.

Our algorithm evaluates the expression to determine whether the current priority combination of appliances are still valid or not, $[|\overline{y}_t - \sum_{n=1}^{K} \hat{y}_t^{(n)}| \leq \delta]$, where K is the set of appliances present in the priority combination and δ is the variation threshold. The variation threshold parameter minimizes the effect of small fluctuations in aggregate power. Since these fluctuations vary for different appliances based on their power rating, the δ value needs to be adaptive. However, when the difference between current priority combination and aggregate consumption is greater than δ, we find the new state combination of appliances that match the aggregated energy consumption.

This work distinguishes the appliances in the household into three categories: (i) *Personal appliances*: Personal laptop, hair dryer, smartphones, etc. (ii) *Shared appliances*: Television (TV), kitchen utilities, boiler, microwave, etc. (iii) *Baseline appliances*: appliances that are always ON – modems, routers, refrigerator, etc.

3.2 User Monitoring

User monitoring component is activated upon the reception of an event trigger from appliance monitoring component. User monitoring determines the current location and activities performed by all the occupants. Smartphones are used for indoor room level localization and smartwatches are used for micro-activity recognition.

Location monitoring

This work focuses on smartphones/watches to determine indoor location of occupants due to the following reasons: (i) smartphone/watch is personally associated and carried by a user, (ii) change in sensor information such as accelerometer can be used to detect user movements and (iii) localization techniques can use WiFi and/or Bluetooth radios to identify user location. The event trigger from the appliance monitoring component initiates the data collection for indoor localization. The data stream includes a scan of visible WiFi access points (APs) and their Received Signal Strength (RSS) along with the timestamp. The list of APs indicates the access points from the neighboring houses. To save battery and also to derive accurate location, a scan is performed only upon the detection of a user movement (i.e., change in accelerometer data or step detection).

3.2.1 Naïve Bayesian classifier for localization. Classification techniques such as Bayesian, Support Vector Machines, K-nearest neighbor and decision trees, have been proposed in the literature to derive room level occupancy using RSS information. Our localization algorithm is based on Bayesian classification technique[5] and has two phases *viz.,* training and testing phase. During the training phase, data is collected at each room to build a classifier model. This phase is also called the fingerprinting stage, where data from WiFi scans are used to learn the available APs and their RSS at different locations.

Feature extraction: The collected data from WiFi scan is then used to derive features for the classifier model. Feature vectors are derived by using multiple scans performed. In this work, we use four WiFi scans to derive feature vectors such as max, min, mean, standard deviation of the signal strength for each available AP. Feature vector $\mathbf{l}$ for k access points is represented as,

$$\begin{aligned} l_t =&< rss_t^{max}(1), rss_t^{min}(1), rss_t^{mean}(1), rss_t^{std}(1), \ldots, \\ & rss_t^{max}(k), rss_t^{min}(k), rss_t^{mean}(k), rss_t^{std}(k) > \end{aligned} \tag{2}$$

Building classifier model: Feature vectors obtained are provided as input to the classifier algorithm to derive the class labels. The classifier model generates a probability distribution function (PDF), which is further used in testing phase to determine the class label (room location).

Activity monitoring

Wearable devices provide an opportunity to identify precisely the micro-activities performed by the occupant such as opening microwave door, opening the refrigerator door, using a laptop, switching on/off an appliance. Smartwatch of an occupant is used to determine the activity performed. The different sensors used for acquiring relevant activity information include an accelerometer, gyroscope, magnetometer, tilt, rotation and linear acceleration. The activity feature vector, $\mathbf{a}$, includes data from both time and frequency domain. Features considered are mean, standard deviation, variance from x, y, z axes along with magnitude, fundamental frequency, zero crossing and step counter.

The event trigger initiates the data collection for activity monitoring. Similar to location monitoring, the activity monitoring also has training and testing phase. During the training phase, multiple samples of activity feature are collected for each micro-activity. The activity feature vectors are then used by the classifier to model and determine the class label associated with each activity feature. The class label indicates the micro-activity performed by the occupant. In the testing phase, each activity feature is evaluated by the activity model to determine the micro-activity performed. Location and activity monitoring determine *where* the occupant is and *what* activities are performed by the occupant when an event trigger is received.

[4]OKSOLAR. [Online] http://www.oksolar.com/technical/consumption.html

[5]Note that other classification techniques can also be used to derive location information.

3.3 Appliance-User Modeling

The appliance-user modeling studies the user association with the appliance. The objective of the appliance-user modeling is to determine the occupants that are currently associated with the appliance being used. If the current appliance being used is personal, then it can be assigned to the relevant user. However, as seen in Fig. 1 some appliances are shared by all the occupants at different time periods (e.g., Microwave) and some appliances are used by all the occupants at the same time period (e.g., TV). Hence it is important to not just determine where the occupant currently is but also to determine the activities performed by the occupant. PEAT utilizes location and activity information to determine the user association with the appliance.

To study the association between the appliances and users we first determine the event type. PEAT distinguishes the event triggers from appliance monitoring into two *viz.,* (i) unambiguous events and (ii) ambiguous events. *Unambiguous events* are those when there is a total certainty that a single occupant is using the appliance. These events occur when there is only one occupant in the household at that time or when a personal appliance is used. By filtering the occupants who are outside the household, the model determines the occupant currently using the appliance. Unambiguous events can help to determine characteristics of occupants such as, which appliances are used by only one occupant at an instance? Which appliances are commonly used by a specific occupant? Furthermore, for an *ambiguous event* there is more than one occupant associated with the usage of an appliance. This is typically the case in multi-occupant households, where cooking, watching TV, and using lights are usually group activities.

In the case of an ambiguous event, PEAT first determines the location of the occupant by evaluating the location feature vector. All the occupants whose predicted location is different from the location where the appliance event occurred are discarded. Furthermore, for all occupants who are in the same location as the appliance is used, the system identifies the occupants who are associated with a smartwatch. The activity information from these occupants is evaluated to identify the micro-activities performed. The model determines if the activity performed resembles the activity related to the appliance usage and associates the occupant accordingly. For example, if the activity of an occupant inferred is "opening microwave door" and the appliance event is from microwave, then PEAT assigns the usage of the microwave to that occupant. Energy is apportioned with a high certainty to the occupant whose activity matches with that of the appliance being used. *Note that, not all users in the household may have a smartwatch/phone.* Hence, there could be an ambiguous event that is not being resolved. Most of the energy apportioning systems, hitherto, divide evenly the energy consumption if they cannot resolve the ambiguous events. To this end, PEAT utilizes several heuristics that relies on historical data of the occupant's association with the appliances.
(i) *Number of times used* ($\mathbf{H}_1$) assigns higher association probability to the occupant who has used the appliance more number of times.
(ii) *Recently used* ($\mathbf{H}_2$) assigns higher association probability to the occupant who recently used the appliance.
(iii) *Average usage duration* ($\mathbf{H}_3$) assigns higher association probability to the occupant whose average appliance usage duration matches the current appliance usage duration.

These heuristics are used to determine the percentage of energy that needs to be distributed among the occupants when an ambiguous event is not resolved. For example, if the appliance used is in a Bedroom, H_1 assigns a higher probability to the user of that bedroom. Similarly, if the appliance used is Refrigerator, then H_2 assigns a higher probability to the occupant who recently used. This is generally the case in the kitchen. If an Occupant-A watches TV for approximately an hour and Occupant-B watches TV generally for 30 m, then the heuristic H_3 assigns a higher probability to Occupant-A when the current TV usage exceeds 30 m. Characteristics such as a number of times an appliance is used, its average usage duration, are learned over time by analyzing the energy consumption pattern and location information of occupants. Finally, after applying all the heuristics, the event is assigned to a single occupant or group of occupants who are more likely to have used the appliance. If the association probability to one or more occupants has similar values, then the event is assigned to all those occupants and the energy is equally apportioned. The association probabilities and usage characteristics of occupants are stored in Raspberry Pi for deriving per-occupant statistics and to adapt the heuristics over time.

3.4 Online Evaluation

This component evaluates the energy to be apportioned to each occupant in the household in real-time. The evaluation starts when the "OFF" event trigger is obtained from the appliance. The event is then classified to be either ambiguous or unambiguous based on the location information and appliance under consideration. PEAT then evaluates the location and activity information obtained from the user monitoring component. The location accuracy may be inaccurate in some scenarios due to misclassification or the user may not have carried his/her phone. To overcome this PEAT applies a simple location correction mechanism. From the metadata collected, we know the location of each appliance in the household. If there is only one occupant and his location is other than the location of the appliance being used, we then use the appliance location as the corrected location. This corrected location information is then used by the appliance-user modeling component. The association probability derived from the appliance-user modeling is used to apportion the energy among occupants. This information is further sent to all the occupants with individual and shared energy consumption details.

4 EXPERIMENTAL SETUP

4.1 Deployment details

To evaluate PEAT in real-world, the complete system was deployed in two multi-occupant settings *viz.,* student housing and a residential household. The *student house* is a two-bedroom apartment with *four* locations *viz.,* Kitchen, Living room, Bedroom 1 and Bedroom 2 as illustrated in Fig. 3(a). All locations apart from the bedrooms are shared by the occupants. The appliances include microwave, refrigerator, grill, coffee machine, laptops and television (TV). Three

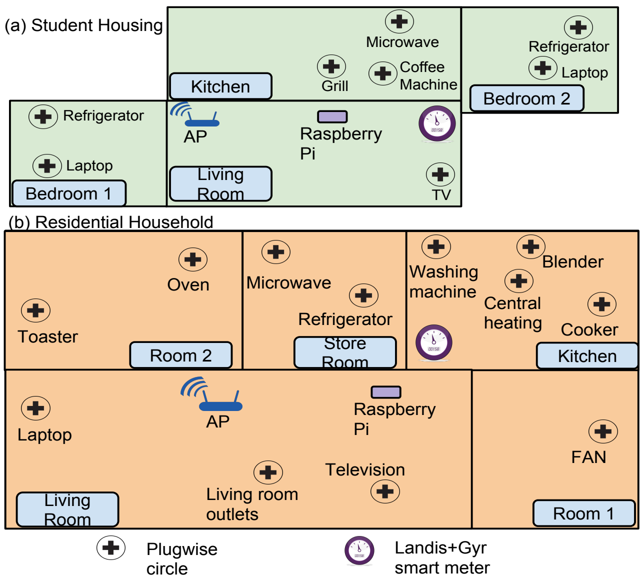

Figure 3: Floor plan of student housing and residential household.

occupants were present in the student house during our experimentation.

The *residential household* contains 12 appliances spread across 5 rooms as shown in Fig. 3(b). Two occupants were present in the household during our experimentation.

4.2 Sensing infrastructure and data collection

Our deployment consists of several sensors measuring electricity, occupancy and activities of occupants. The sensors were carefully installed to avoid any inconvenience to the occupants.

Electricity monitoring: We used off-the-shelf sensors to monitor energy consumption at 1 Hz sampling frequency.

(i) *Mains level*: We installed a smart electricity meter from Landis+Gyr E350 to measure the aggregate energy consumption information. The data from the smart meter was retrieved using Plugwise Smile[6].

(ii) *Appliance level*: We used smart plugs from Plugwise circle[7] to collect appliance level energy consumption data.

The plugs communicate via Zigbee protocol by forming a mesh network. We use an open source library python-plugwise to query the data from the plugs at 1 Hz frequency. A Raspberry Pi was deployed locally to generate periodic queries and to store the data.

Occupancy monitoring: In our deployment, we scan for visible WiFi access points using smartphone/smartwatch of occupants for indoor localization. Furthermore, the smartphones carried by the occupants included: Samsung Galaxy III, Nexus 7 and LG Nexus 5. An Android application was developed to scan WiFi APs upon reception of an event. This data is further used by the online classifier on the phone to determine the indoor room level location.

Activity monitoring: Smartwatches like Samsung galaxy gear 2 and Moto 360 was used to determine the micro-activities performed.

[6]Smile:https://www.plugwise.com/smile-p1
[7]Circle:https://www.plugwise.com/circle

Table 2: Metadata: association of appliance and location.

Student housing			Residential household	
Appliance	**ID**	**Location**	**Appliance**	**Location**
Microwave	E_1	Kitchen	TV	Living Room
Laptop	E_2	Bedroom 1	Fan	Room 1
Refrigerator	E_3	Bedroom 1	Laptop	Living Room
Refrigerator	E_4	Bedroom 2	Microwave	Store room
Laptop	E_5	Bedroom 1	Oven	Room 2
TV	E_6	Living Room	Toaster	Room 2
Grill	E_7	Kitchen	Central heating	Kitchen
Coffee Machine	E_8	Kitchen	Living room outlets	Living Room
			Blender	Kitchen
			Washing machine	Kitchen
			Cooker	Kitchen
			Refrigerator	Store Room

Sensors such as accelerometer, gyroscope, magnetometer, tilt, rotation, and linear acceleration were used for activity recognition.

Household metadata: PEAT utilizes the list of appliances and its location in the household as the metadata. This is a one-time activity and can be obtained during the setup. Table. 2 shows the association of the appliance and location in both the settings.

Ground Truth: To validate the results from PEAT, the ground truth about the use of an appliance is required. Hence, we deployed NFC tags to collect this information. Each tag is pre-programmed with the appliance name and location. Upon the initialization of an NFC tag [36], an event is logged into the system with the occupant ID. We use this information only for comparing the results derived from PEAT.

4.3 System architecture

The system architecture consists of several clients (occupants' devices) communicating with a server (local raspberry PI). *Server-side* includes a Raspberry PI with WiFi connectivity that acts as the local server in each household. Raspberry PI receives the energy consumption data from the smart meter. In our setup, we used Plugwise Smile-P1 to retrieve the data from the smart meter and send it to the Raspberry PI via WiFi. Raspberry PI runs the proposed ModCO energy disaggregation algorithm to derive fine-grained appliance usage information. Upon detection of appliance ON event, Raspberry PI sends out a push notification (trigger) to all the clients (occupants devices i.e., smartphones/watches) to start scanning (i) WiFi RSSI samples for indoor localization on smartphones and (ii) data collection of inertial sensors on smartwatches of users. Further, when an appliance OFF event is detected, Raspberry PI sends out another push notification to stop the data collection at the client devices. On the *client-side*, an application is developed for smartphones and watches of users. There could be multiple occupants in a household and hence during the initial phase each occupant is assigned a unique ID along with their devices. To save energy on the client devices, we do not start data collection until a trigger is received from the local server. Upon reception of the trigger on smartphones, the application starts scanning for WiFi RSSI signals and then sends out a notification to the smartwatch associated for collecting activity related information. Further, when an OFF event trigger is received, the application utilizes an online version of Bayesian classification to derive the room level occupancy and the micro-activity performed.

Further, Raspberry PI receives the inferred location and micro-activity performed for the ON-OFF event of each appliance. This information is used to develop an appliance-user modeling and for apportioning energy to individual occupants. The proposed system architecture is distributed, wherein, localization and activity recognition is performed on smartphones/smartwatches, and energy disaggregation and apportioning are performed on the Raspberry PI. Note that all the devices, Plugwise Smile-P1, Raspberry PI, and smartphones, are connected to the same access point in the household.

5 EVALUATION

We evaluated each component of the toolkit extensively in both student housing and residential household. Appliance monitoring component was evaluated with both original CO and the proposed modified CO. We evaluated location and activity feature vectors across three online classifiers *viz.,* Decision trees (J48) [29], Naïve Bayesian (NB) [31] and K-Nearest Neighbors (KNN) [30]. Finally, we studied the trade-off between energy apportioning accuracy and the number of devices (smartphones/smartwatches). The following methods were employed to derive energy apportioning accuracy:
M$_1$: One user with the smartwatch, all other users with the smartphones and with heuristics H$_1$, H$_2$ and H$_3$.
M$_2$: No smartwatch, all users with the smartphones and with heuristics H$_1$, H$_2$ and H$_3$.
M$_3$: All users with the smartwatches, smartphones and with heuristics H$_1$, H$_2$ and H$_3$.
M$_4$: One user with a smartwatch, all other users with the smartphones and no heuristics.
M$_5$: No smartwatch, all users with the smartphones and no heuristics.
Metrics: Several accuracy metrics are considered here to evaluate the components of the toolkit. Different metrics considered for appliance monitoring are given below.

Fraction of total energy assigned correctly (FTE): It measures the fraction of energy correctly assigned to an appliance and is one of the common accuracy metrics for NILM algorithms [10, 21]. It is defined as,

$$FTE = \sum_n min\left(\frac{\sum_n y_t^{(n)}}{\sum_{n,t} y_t^{(n)}}, \frac{\sum_n \hat{y}_t^{(n)}}{\sum_{n,t} \hat{y}_t^{(n)}}\right), \tag{3}$$

where $n \in \{1, .., N\}$, N is the total number of appliances, $t \in \{1, .., T\}$ and T is the total time period considered.

Total disaggregation error (T_e): Total disaggregation error is the difference between the total energy consumed by all the appliances and the actual energy consumed by the appliances, normalized by the total energy consumed. It is given by,

$$T_e = \frac{\sum_{n,t} |y_t^{(n)} - \hat{y}_t^{(n)}|}{\sum_{n,t} y_t^{(n)}} \tag{4}$$

Number of appliances identified correctly (J_a): Jaccard similarity coefficient is used to measure the similarity between the predicted set of appliances (J_a^p) and the actual set of appliances (J_a^a) used over a period of time. J_a measures the percentage of appliances correctly identified by the disaggregation algorithm. It is given by,

$$J_a = \frac{|J_a^p \cap J_a^a|}{|J_a^p \cup J_a^a|} \tag{5}$$

Number of appliance states identified correctly (J_s): It measures the similarity between the predicted set of appliance states (J_s^p) and the actual set of appliance states (J_s^a). It is given by,

$$J_s = \frac{|J_s^p \cap J_s^a|}{|J_s^p \cup J_s^a|} \tag{6}$$

We now describe the set of metrics used to evaluate the classifier models obtained for location and activity.

Precision: It is the ratio of number of correctly identified instances over total number of identified instances. Let t_p and f_p indicate the true positives and false positives respectively and precision is defined as,

$$precision = \frac{t_p}{t_p + f_p} \tag{7}$$

Recall: It is the ratio of number of correctly identified instances over total number of instances. It is given by,

$$recall = \frac{t_p}{t_p + f_n}, \tag{8}$$

where f_n represents the number of false negatives.

F$_1$ Score: It is a measure of accuracy and is defined as the harmonic mean of precision and recall.

$$F_1 = 2.\frac{precision \cdot recall}{precision + recall} \tag{9}$$

Energy apportioning accuracy (E_a): Energy apportioning accuracy is the ratio of estimated energy utilized by an occupant and the actual energy utilized by that occupant. It measures the total percentage of energy correctly apportioned to an occupant and is given by,

$$E_a = \left(\frac{\text{Estimated energy per occupant}}{\text{Actual energy per occupant}}\right) \times 100 \tag{10}$$

6 RESULTS

In this section, we present our experimental results in determining state change of an appliance, room level occupancy, and activities performed by the occupants. Furthermore, we show the energy apportioning accuracy across different real-world multi-occupant settings and its trade-off with respect to the number of devices used.

6.1 Appliance detection accuracy

PEAT employs modified CO to determine the state of the appliances in real-time. Accurate energy disaggregation is a critical component for unambiguous energy apportioning. To ensure a fair comparison, both original and ModCO utilize the same appliance model from the crowd-sourced database as described in Section 3.1. Furthermore, comparison with other NILM algorithms (FHMM [9]) requires additional training such as prior probability and state transition matrix. Hence we restrict the comparison of proposed ModCO with the original CO algorithm.

Fig. 4 shows the disaggregation performance of CO and modified CO across *four* accuracy metrics in the residential household. We used over 2 months of aggregated energy consumption data of the household. FTE, J_a and J_s can vary between 0 and 1, and T_e can take any non-negative value. ModCO assigns (FTE) more than 85% of the aggregate energy accurately. Furthermore, around 75% of state changes (J_s) are estimated correctly as compared to 35% by original CO. Similarly, when a student housing was considered 90% of total

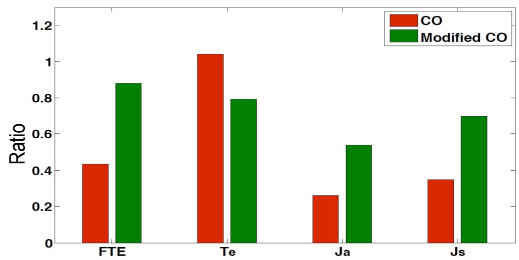

Figure 4: Disaggregation performance of CO and modified CO in residential household.

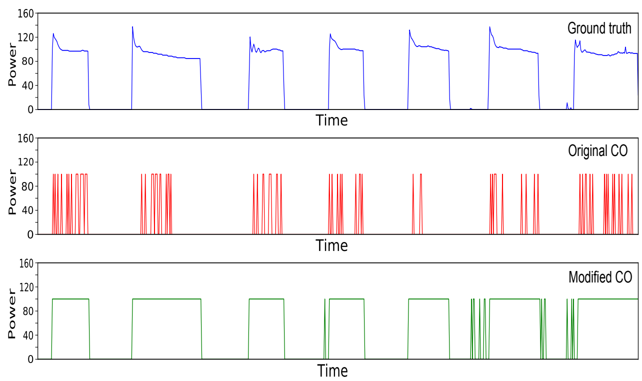

Figure 5: Original and disaggregated energy profile of refrigerator using CO and modified CO.

energy was accurately assigned and 87% of the states were identified correctly. The increase in FTE and appliance detection in student housing compared to the residential household is mainly due to the reduced number of appliances considered. In both the settings, T_e in modified CO is much lower than original CO, indicating better disaggregation performance.

In general, the improvement in disaggregation accuracy for ModCO is due to the fact that the predicted set of appliances does not vary significantly for consecutive time periods. Fig. 5 shows the estimated energy by original CO and modified CO as compared to the ground truth. It can be clearly seen that during consecutive predictions, CO is very sensitive to changes in aggregated energy consumption leading to the selection of the wrong state. However, ModCO overcomes this with the help of the priority combination and the δ parameter. Moreover, this results in significant reduction of associated computational complexity. Our ModCO takes approximately 0.12 s and 0.05 s to determine the state of the appliances in the residential household and student housing, respectively.

6.2 Indoor location accuracy

An Android application was installed on smartphones to scan for visible WiFi APs when an event trigger is received. Feature vectors were computed as described in Section 3.2. During our experimentation, we considered two approaches for building the classifier model *viz.,* supervised and unsupervised.

Supervised method requires a training phase where RSS values at each location are collected and labeled. During the testing phase, each feature vector is evaluated with the classifier model obtained in the training stage to derive the class label (room level location). *Unsupervised method* does not know the class labels *a priori* and learns the label of the location based on occupancy and the appliance metadata. For example, when only one occupant is present in the household and if the appliance trigger was from "Coffee Machine" then the location of the appliance (i.e., Kitchen) can be obtained from Table. 2. Consequently, the algorithm learns this label and assigns it to the current location of the occupant. Furthermore, this iterative approach continues until all the class labels are determined. However, this method has several drawbacks such as works when only one user is present in the household and a longer delay in developing accurate location models. Recently, several algorithms such as Zee [27] and EchoTag [28] employ crowd-sourced data collection to eliminate the tedious training phase. These approaches could be also used in PEAT. To show the effectiveness of PEAT, we use a standard online classification model.

Training data: We developed an Android app to collect WiFi information in each room and assign a label. The user collected WiFi fingerprinting data for each room once every two hours. WiFi fingerprinting data was collected over a period of 3 weeks, which was then used to build the classification model.

We employed Naïve Bayesian (NB) classifier model [31] to derive class labels for each new feature vector obtained. Fig. 6(a) shows the precision and recall for each location in the student housing. High values of precision and recall at each location indicate the good performance of the classifier model. Moreover, F_1 score of 84% was achieved for room level localization using NB. Furthermore, we compared the classification results with two other well-known classifiers *viz.,* J48 [29] (Decision trees) and KNN [30] with 10-fold stratified cross validation. Fig. 6(c) shows that NB performs much better than J48 and KNN, with KNN having the least classification accuracy. Finally, F_1 score of 78% was achieved using NB for room level localization in the residential household. NB does not over-fit the data as compared to J48 and KNN. Moreover, NB outperforms J48 and KNN even with partial data collected from occupantsŠ smartphones.

6.3 Activity detection accuracy

In this work we considered *six* micro-activities *viz.,* (i) microwave usage (A_1), (ii) laptop usage (A_2), (iii) refrigerator usage (A_3), (iv) TV usage (A_4), (v) grill usage (A_5) and (vi) coffee machine usage (A_6). When there is an "ON" event from appliance monitoring, activity information is collected from occupants using the smartwatch as described in Section 3.2.

Training data: Each micro-activity was performed by the user wearing a smart watch. We collected around 40-50 instances of labelled data for each micro-activity. The ground truth was collected by manually annotation and using NFC tags.

In the training phase, each activity is labeled and the features in both time and frequency domain are collected. Through exhaustive experiments, we found that 6 s of samples at 100 Hz sampling frequency as optimal to detect the activities accurately. Note that a large sampling duration may include additional information, which may not be relevant and having a small sampling duration may not capture the relevant features of an activity. Hence, identifying

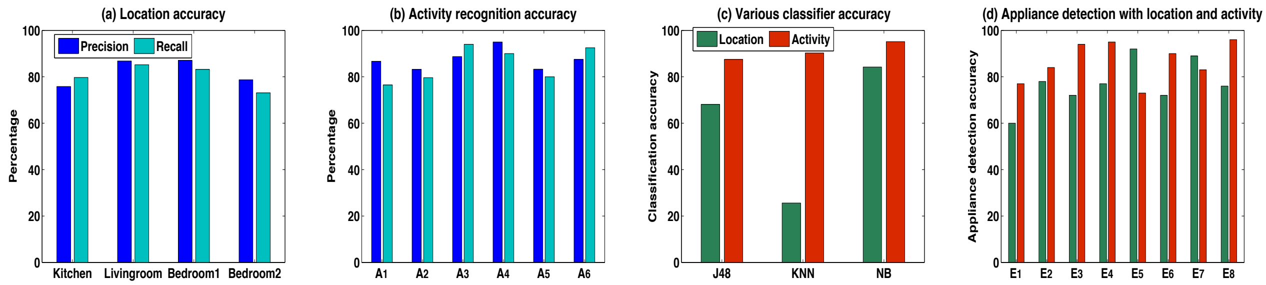

Figure 6: Evaluation of user monitoring component in student housing.

Table 3: Percentage of energy correctly apportioned for different methods in student housing.

Methods	$\mathbf{M_1}$	$\mathbf{M_2}$	$\mathbf{M_3}$	$\mathbf{M_4}$	$\mathbf{M_5}$
Energy apportioning accuracy	92.6	90.7	95.4	87.1	80.6

the right sampling duration is crucial in determining the micro-activities performed.

Similar to location monitoring we evaluated the activity feature vector across three classifiers (see Fig. 6(c)). NB has a classification accuracy of 95% where as J48 and KNN have around 87.5% and 90%, respectively. Fig. 6(b) shows the precision and recall for each micro-activity. It can be seen that the precision of all the activities are higher than 85% and the overall accuracy of identifying the activities is around 95%. Even though the activities A_1 and A_3 are quite similar, the classifier was still able to identify them correctly. This is attributed to the correct identification of the sampling duration.

Furthermore, we conducted several experiments to understand the effectiveness of the location and activity features with respect to appliance detection. The accuracy of detecting an appliance usage with either location or activity features is shown in Fig. 6(d). x-axis indicates the various appliances in the student housing (see Table. 2) and y-axis represents the appliance detection accuracy. In general, the activity features can determine the associated appliances more accurately than location features. This is due to the identification of micro-activities using the smartwatches. It can be seen that for some appliances such as Refrigerator and Grill, location features have higher accuracy than activity features. This is attributed to the placement of appliances in different rooms and their distinctive consumption profile. Moreover, for both location and activity information, the average accuracy of identifying the associated appliance is around 82%. This information can also be used with ModCO to improve the accuracy of appliance detection.

6.4 Energy apportioning accuracy

We considered over 2 months of data to evaluate PEAT in both student housing and residential household. All the occupants were equipped with their personal smartphone and smartwatch. Furthermore, the proposed toolkit can also be applied to other shared spaces such as office environments. Note that the level of apportioning

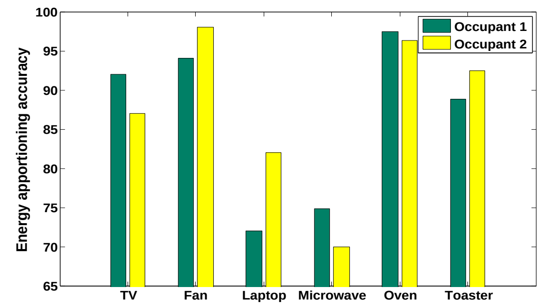

Figure 7: Energy apportioned correctly for each occupant in the residential household.

required in these spaces varies. In student housing and residential household, PEAT apportions energy of shared and personal appliances. The baseline appliances such as refrigerator, routers, air conditioning and central heating are not apportioned to individual occupant. Moreover, the toolkit can be extended to support apportioning of baseline appliances with additional training and micro-activity recognition (e.g., open the refrigerator door and take/keep items).

To study the efficacy of PEAT, we evaluated energy apportioning accuracy for varying number of devices as mentioned in Section 5. Table. 3 shows the total percentage of energy correctly apportioned for a week in student housing with varying number of user devices. When there is one occupant with smartwatch and other occupants have only their smartphones i.e., method M_1 with heuristics, PEAT achieves around 92% apportioning accuracy per-occupant. This accuracy reduces to 87% when no heuristics are considered i.e., M_4. Furthermore, when all the occupants have only their smartphone but no smartwatches, PEAT still achieves 90.7% accuracy with heuristics (M_2) and it is 80% when no heuristics are applied (M_5). Finally, if all the occupants had both smartphone and smartwatch, then the energy apportioning accuracy achieved by PEAT is around 95% (M_3). Clearly, having more user devices increase the apportioning accuracy.

We now illustrate the energy apportioning accuracy for the method M_1 among all occupants on the per-appliance basis. Table. 4 shows appliance level energy that was correctly apportioned to all the occupants in the student housing. It can be seen that when an

Table 4: Percentage of energy correctly apportioned for each appliance among all the occupants.

Appliance ID	E_1	E_2	E_3	E_4	E_5	E_6	E_7	E_8
Apportioning accuracy	66.6	100.0	88.9	88.9	100.0	88.9	80.7	77.8

event was associated with a personal appliance such as Laptops, the apportioning accuracy was 100%. However, with the shared appliances the apportioning accuracy reduces to 80% on average. The average energy apportioning accuracy for all the appliances is close to 90% and per-occupant average apportioning accuracy was around 92% in the student housing.

Fig. 7 shows the energy apportioned to each occupant across appliances in the residential setting when only smartphones were used (M_2). PEAT achieves on average 87% apportioning accuracy for all occupants when only location information was used in a residential household. The apportioning accuracy increases to 92% when all the occupants in the household have their smartphones and smartwatches. Furthermore, for baseline appliance “refrigerator" an apportioning accuracy of 88.9% is obtained for both the occupants.

6.5 Discussions

In this section, we first highlight the design challenges solved in PEAT and then discuss some of the limitations with possible solutions to overcome it.

(i) Apportioning: We proposed a novel modeling of appliance-user relationship to determine which occupant used which appliance and where in the household. We presented three heuristics to further enhance the apportioning accuracy during ambiguous events.

(ii) Scalability: We presented a distributed system architecture with the local embedded server (Raspberry PI) for processing and storage of data in the household. Energy disaggregation, localization, and activity monitoring run on different occupant's devices. The local server determines when to trigger various devices in the household to start and stop data collection, making the system efficient and robust.

(iii) Real-time: We presented our ModCO disaggregation algorithm that runs on an embedded device in real-time. Further, a standard Bayesian classification algorithm was implemented on occupants' smartphones to determine user location/activity in real-time. Energy apportioning algorithm runs on the local server and provides real-time usage details to each occupant.

(iv) End-to-end system deployment: We extensively evaluated PEAT in two multi-occupant settings with varying number of occupant devices. PEAT system setup is non-intrusive, privacy-preserving, easily configurable and deployable.

While PEAT takes the first few steps towards effective user-level energy disaggregation, there are some limitations:

(i) Resolution of apportioning: The level of apportioning may vary depending on the environment, for example, in residential settings it may be more useful to apportion top energy consuming appliances than all household appliances. In shared spaces such as student housing, shared appliance usage such as TV, washing machine, microwave, and oven needs to be apportioned to raise awareness among occupants. In PEAT, users can select appliances for which energy needs to be apportioned to individual occupants during the initial setup.

(ii) HVAC apportioning: HVAC usage is the primary energy guzzler in households. Hitherto, HVAC consumption was equally shared among all occupants. PEAT with the help of energy disaggregation can identify when HVAC was turned on/off and using location monitoring can determine when and where each occupant are in the household. PEAT attributes the HVAC consumption to the users depending on the time they spent in the household. However, this works only for centralized HVAC systems and currently, there is no way to determine individual room HVAC consumption without any additional sensors.

(iii) Shared appliances: PEAT with the help of activity and location monitoring can determine which occupant is using the appliance when multiple occupants are in the same location. However, when an occupant is not carrying his device (smartphone/smartwatch) the event cannot be resolved to an occupant accurately. While we propose three heuristics to apportion energy in such cases, further enhancement of appliance-user modeling is required for accurate apportioning.

(iv) Moveable appliances: PEAT requires location of appliances known *a priori*. Hence it cannot accurately apportion energy usage of moveable appliances such as hair dryer when there are multiple occupants present in the room.

(v) Localization and activity monitoring algorithms: PEAT achieves 90% apportioning accuracy by using standard classification algorithms. We believe PEAT can be more effective and robust by incorporating other crowd-sourced training free algorithms [27, 28].

(vi) Dependency on training data: PEAT is designed as a supervised system where training data related to micro-activities, fingerprinting WiFi signals for localization and household appliance details are collected apriori. This restricts the usage of PEAT on an entirely new household. However, recent algorithms for activity recognition [33–35] and localization [27, 28] overcome the bottleneck of training data. Since the design of PEAT is modular, researchers can use the newer models to reduce training effort and improve accuracy.

(vii) Deployment challenges: As PEAT relies on appliance usage, localization and micro-activities data from all the occupants, collecting and modeling such data from different sources is a challenging task. One key challenge is how to collect ground truth information? - To this end, we deployed NFC tags to detect appliance usage and video cameras to localize occupants and determine their activities. Further, we manually annotated this data to derive ground truth. Another challenge was to deploy appliance-level energy monitoring sensors and NFC tags in a non-intrusive way so that the occupant behavior is unaltered. Recently there have been open datasets [37, 38] that collect smartphone/smartwatch data from multiple occupants in diverse situations. PEAT can ingest this data for robust modeling of user activities and location.

7 CONCLUSIONS AND FUTURE WORK

We proposed a novel Personalized Energy Apportioning Toolkit (PEAT) to accurately apportion energy amongst occupants in shared spaces. Inferring energy footprints of occupants with minimal user intervention and no additional sensor deployment is a challenging task. We showed that PEAT can accurately determine *which* occupant performed *what* activity and *where* in the household. Furthermore, it combines online techniques with minimal training for accurate energy apportioning. We proposed several accuracy metrics to study the performance of each component of PEAT. We specifically deployed the system in two multi-occupant settings – *viz.,* a student housing and a residential household – to evaluate PEAT in the real-world settings. With only the location information, energy apportioning accuracy of 87% and 92% was achieved for all the occupants in the residential household and student housing, respectively. The apportioning accuracy increases to 92% and 95% when both location and activity information was available in the residential household and student housing. We demonstrated that PEAT is highly scalable and privacy-preserving since privacy-sensitive data of occupants are stored and processed locally at the households.

Next we envisage following three directions: (i) extending the toolkit to support data from other publicly available datasets for energy apportioning to increase the performance; (ii) using personalized energy consumption information to enhance existing energy management systems; and (iii) extending PEAT to other shared resources such as, gas and water in shared spaces.

ACKNOWLEDGMENTS

SCOTT http://www.scott-project.eu has received funding from the Electronic Component Systems for European Leadership Joint Undertaking, grant agreement No 737422. This joint undertaking is supported from the EU Horizon 2020 research and innovation program and Austria, Spain, Finland, Ireland, Sweden, Germany, Poland, Portugal, Netherlands, Belgium, Norway.

REFERENCES

[1] Energy outlook 2010. Energy Information Administration. [Online] http://www.eia.doe.gov/oiaf/ieo/index.html, 2010.
[2] M. Hazas, A. Friday, J. Scott, “Look Back before Leaping Forward: Four Decades of Domestic Energy Inquiry,” *In Proc. of Pervasive Computing,* 2011.
[3] A.G. Ruzzelli, C. Nicolas, A. Schoofs, G.M.P. O'Hare, “Real-Time Recognition and Profiling of Appliances through a Single Electricity Sensor,” *In Proc. of IEEE SECON,* 2010.
[4] S. Barker, S. Kalra, D. Irwin, and P. Shenoy, “PowerPlay: creating virtual power meters through online load tracking,” *In Proc. of ACM BuildSys,* 2014.
[5] R. Fontugne, J. Ortiz, N. Tremblay, P. Borgnat, P. Flandrin, K. Fukuda, D. Culler, H. Esaki, “Strip, bind, and search: a method for identifying abnormal energy consumption in buildings,” *In Proc. of Information processing in sensor networks (IPSN),* 2013.
[6] V. Srinivasan, J. Stankovic, K. Whitehouse, “FixtureFinder: discovering the existence of electrical and water fixtures,” *In Proc. of Information processing in sensor networks (IPSN),* 2013.
[7] B. Ho, H. C. Kao, N. Chen, C. You, H. Chu, and M. Chen, “HeatProbe: a thermal-based power meter for accounting disaggregated electricity usage", *In Proc. of Ubiquitous computing UbiComp,,* 2011.
[8] S. Darby, “The effectiveness of feedback on energy consumption,” *A review for DEFRA of the literature on metering, billing and direct displays,* 2006.
[9] S. N. Akshay Uttama Nambi, T G. Papaioannou, D Chakraborty and K Aberer, “Sustainable energy consumption monitoring in residential settings,” *In Proc. of the IEEE INFOCOM,* 2013.
[10] J. Zico Kolter and M. J. Johnson, “REDD: A public data set for energy disaggregation research,” *In Proc. of the SustKDD Workshop on Data Mining Applications in Sustainability,* 2011.
[11] Y. Guo, M. Jones, B. Cowan, and R. Beale, “Take it personally: personal accountability and energy consumption in domestic households,” *In Proc. CHI '13 on Human Factors in Computing Systems,* 2013.
[12] D. Bonino, F. Corno, and L. D. Russis, “Home energy consumption feedback: A user survey,” *In Energy and Buildings,* vol. 47, pp. 383-393, 2012.
[13] U. Gneezy, E. Haruvy, and H. Yafe, “The inefficiency of splitting the bill,” *In The Economic Journal,* vol. 114(495), pp. 265-280, 2004.
[14] S. Hay and A. C. Rice, “The case for apportionment,” *In Proc. ACM BuildSys,* 2009.
[15] S. Lee, D. Ahn, S. Lee, R. Ha, and H. Cha, “Personalized energy auditor: Estimating personal electricity usage,” *In Proc. of the International Conference on Pervasive Computing and Communications,* 2014.
[16] G. W. Hart, “Nonintrusive appliance load monitoring,” *In Proc. of the IEEE Journal,* vol. 80(12), pp. 1870-1891, 1992.
[17] S. N. A. U. Nambi, “Personalized Energy Services: A Data-Driven Methodology towards Sustainable, Smart Energy Systems", Delft University of Technology, 2017.
[18] Y. Cheng, K. Chen, B. Zhang, C.-J. M. Liang, X. Jiang, and F. Zhao, “Accurate real-time occupant energy-footprinting in commercial buildings,” *In Proc. of ACM BuildSys,* 2012.
[19] M. Saha, S. Thakur, A. Singh, and Y. Agarwal, “EnergyLens:Combining Smartphones with Electricity Meter for Accurate Activity Detection and User Annotation,” *In Proc. of 5th International Conference on Future Energy Systems,* 2014.
[20] J. Ranjan, E. Griffiths, and K. Whitehouse, “Discerning electrical and water usage by individuals in homes,” *In Proc. of ACM BuildSys,* 2014.
[21] N. Batra, J. Kelly, O. Parson, H. Dutta, W. Knottenbelt, A. Rogers, A. Singh, M. Srivastava, “NILMTK: An Open Source Toolkit for Non-intrusive Load Monitoring,” *In Proc. of 5th International Conference on Future Energy Systems (ACM e-Energy),* 2014.
[22] S. N. A. U. Nambi, A. Reyes Lua and R. Prasad, “LocED: Location-aware energy Disaggregation Framework,” *In Proc. of ACM BuildSys,* 2015.
[23] V. L. Erickson, S. Achleitner, A. E. Cerpa, “POEM: power-efficient occupancy-based energy management system,” *In Proc. of international conference on Information processing in sensor networks (IPSN),* 2013.
[24] J. Ranjan, Y. Yao, and K. Whitehouse, “An RF doormat for tracking people's room locations,” *In Proc. of the ACM International Joint conference on Pervasive and Ubiquitous Computing,* 2013.
[25] M. Azizyan, I. Constandache, and R. Choudhury, “SurroundSense: mobile phone localization via ambience fingerprinting,” *In Proc. of the 15th annual international conference on Mobile computing and networking (MobiCom),* 2009.
[26] International Data Corporation (IDC) [Online] http://www.idc.com/getdoc.jsp?containerId=prUS25903815
[27] A. Rai, K. Chintalapudi, V. Padmanabhan, and R. Sen, “Zee: zero-effort crowdsourcing for indoor localization", *In Proc. of ACM Mobicom,* 2012.
[28] Y. Tung and K. Shin, “EchoTag: Accurate Infrastructure-Free Indoor Location Tagging with Smartphones", *In Proc. of ACM Mobicom,* 2015.
[29] B. Othman, M. Fauzi, and T. Moh Shan Yau, “Comparison of different classification techniques using WEKA for breast cancer", *In Proc. of 3rd Kuala Lumpur International Conference on Biomedical Engineering, Springer, Berlin,* 2007.
[30] T. Grigorios, and I. Katakis, “Multi-label classification: An overview", *In International Journal of Data Warehousing and Mining* , 2006.
[31] K. Murphy, “Naive bayes classifiers,” *University of British Columbia,* 2006.
[32] P. Wei, X. Chen, J. Vega, S. Xia, R. Chandrasekaran, X. Jiang, “ePrints: A Real-Time and Scalable System for Fair Apportionment and Tracking of Personal Energy Footprints in Commercial Buildings", *In Proc. of 3rd ACM International Conference on Systems for Energy-Efficient Built Environments (BuildSys),* 2017.
[33] J. Wang, Y. Chen, S. Hao, X. Peng, and L. Hu, “Deep learning for sensor-based activity recognition: A survey,” *arXiv preprint arXiv:1707.03502* ,2017.
[34] S. Bhattacharya, and N. D. Lane, “From smart to deep: Robust activity recognition on smartwatches using deep learning,” *In Proc. IEEE International Conference on Pervasive Computing and Communication Workshops (PerCom Workshops),* 2016.
[35] V. Radu, N. D. Lane, S. Bhattacharya, C. Mascolo, M. K. Marina, and F. Kawsar, “Towards multimodal deep learning for activity recognition on mobile devices.", *In Proc. of the International Joint Conference on Pervasive and Ubiquitous Computing: Adjunct ACM,* 2016.
[36] S. N. A. U. Nambi , T.V. Prabhakar, H.S. Jamadagni, et al, “Near Field Communication - Applications and Performance Studies,” *In Proc. of springer CCIS,* 2012.
[37] Y. Vaizman, K. Ellis, and G. Lanckriet, “Recognizing Detailed Human Context In-the-Wild from Smartphones and Smartwatches", *In Proc. of the IEEE Pervasive Computing, vol. 16, no. 4,* 2017.
[38] Y. Vaizman, K. Ellis, G. Lanckriet, and N. Weibel, “ExtraSensory App: Data Collection In-the-Wild with Rich User Interface to Self-Report Behavior", *In Proc. of the 2018 CHI Conference on Human Factors in Computing Systems (CHI 2018), ACM* 2018.