

EnergyLens: Combining Smartphones with Electricity Meter for Accurate Activity Detection and User Annotation

Manaswi Saha[†], Shailja Thakur[†], Amarjeet Singh[†], Yuvraj Agarwal[‡]

[†]Indraprastha Institute of Information Technology, Delhi [‡]Carnegie Mellon University
[†]{manaswis, shailja1275, amarjeet}@iiitd.ac.in, [‡]yuvraj.agarwal@cs.cmu.edu

ABSTRACT

Inferring human activity is of interest for various ubiquitous computing applications, particularly if it can be done using ambient information that can be collected non intrusively. In this paper, we explore human activity inference, in the context of energy consumption within a home, where we define an “activity” as the usage of an electrical appliance, its usage duration and its location. We also explore the dimension of identifying the occupant who performed the activity. Our goal is to answer questions such as “Who is watching TV in the Dining Room and during what times?”. This information is particularly important for scenarios such as the apportionment of energy use to individuals in shared settings for better understanding of occupant’s energy consumption behavioral patterns. Unfortunately, accurate activity inference in realistic settings is challenging, especially when considering ease of deployment. One of the key differences between our work and prior research in this space is that we seek to combine readily available sensor data (i.e. home level electricity meters and sensors on smartphones carried by the occupants) and metadata information (e.g. appliance power ratings and their location) for activity inference.

Our proposed EnergyLens system intelligently fuses electricity meter data with sensors on commodity smartphones – the Wifi radio and the microphone – to infer, with high accuracy, which appliance is being used, when its being used, where its being used in the home, and who is using it. EnergyLens exploits easily available metadata to further improve the detection accuracy. Real world experiments show that EnergyLens significantly improves the inference of energy usage activities (average precision = 75.2%, average recall = 77.8%) as compared to traditional approaches that use the meter data only (average precision = 28.4%, average recall = 22.3%).

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.
e-Energy'14, June 11–13, 2014, Cambridge, United Kingdom.
Copyright is held by the owner/author(s). Publication rights licensed to ACM.
ACM 978-1-4503-2819-7/14/067 ...\$15.00.
<http://dx.doi.org/10.1145/2602044.2602058>.

Categories and Subject Descriptors

C.3 [Special-Purpose and Application-Based Systems]: Real-time and embedded systems

Keywords

Smart Meters; Smartphones; Energy Disaggregation; Activity Detection; User Association

1. INTRODUCTION

Smartphones, over the past few years, have seen unprecedented growth across the world. In many markets, such as the US, they have long since surpassed sales of traditional feature phones. While many factors have contributed to their popularity, the most important perhaps is the increase in device capabilities as supported by higher end components such as multi-core processors, ample memory and storage, and a plethora of embedded sensors. These sensors in turn are used in novel ways by smartphone apps as well as by researchers to build context aware systems [15, 20, 30].

Our work explores the use of these pervasive smartphone platforms for detailed activity inference of individuals within a home setting, which can help occupants make better decisions around their routine daily activities. In particular, we consider the context of activities that lead to energy usage by occupants in a residential setting. This is particularly important since buildings are known to consume significant proportion of energy consumption in both developed and developing countries (45% to 47% of the total energy consumption in the US [5] and India [35] respectively). Smart electricity meters have also been widely deployed with the intention of having complete coverage over the next decade even in the developing countries such as India. Wide scale adoption of smart meters is motivated by several reasons including remote and easy data collection for billing purposes and implementation of new billing practices such as time of day based pricing.

To provide users with finer grained breakdown of their energy usage, there has been much research within the context of direct energy metering [34] or indirect inference using Non-Intrusive Load Monitoring (NILM) techniques [13, 22]. NILM techniques in particular are attractive since they often require a single smart meter and then employ complex machine learning algorithms [14] to disaggregate energy usage at the appliance level. To further improve the accuracy of NILM algorithms, researchers have proposed additional sensors such as EMI detectors [8, 26] and room-level motion and light sensors [27], which are often impractical from an

ease of deployment perspective. To our knowledge, the combination of smartphones and smart electricity meters have not been used before for activity inference in the context of understanding energy usage.

In this paper, we present EnergyLens, a system that leverages two easily available data sources within homes – energy usage data from smart meters and sensor data from smartphones – for accurate activity inference and its annotation to individual home occupants. We define an activity as the usage (*when*) of an electrical appliance (*what*) within a home, as attributed to a specific occupant (*who*), along with the room (*where*) the appliance is located. In other words, our goal with EnergyLens is to answer questions such as “Who is watching the TV in the dining room and during what times?”. Energy consumption information at this granularity is particularly useful to apportion energy usage among occupants in a shared setting as well as understand behavioral patterns for energy usage. Prior work [6] shows, for example, that providing feedback to users about their energy use can help bring about a change in their usage.

Our primary goals in the design of EnergyLens were ease of deployment, widespread applicability and low cost data collection. As a result, while there are numerous sensors available on a modern smartphone we chose to use only WiFi and audio from the microphone since they are available even on low cost devices. Similarly, we only assume the availability of total real power data from a smart meter. While meters that provide additional information such as power factor are available, they are typically more expensive and are not widely deployed in our experience. This information is augmented with metadata information that tags appliances with their power consumption and room location for use by EnergyLens. This metadata was chosen since it is primarily static and requires a one time effort to collect it per home.

EnergyLens therefore combines WiFi based localization and audio-based appliance detection with NILM for appliance classification using meter data. Importantly, we show that even simple algorithms for each of the components – localization, audio classification, and NILM – when used in combination yield good overall accuracy for user activity detection. To show the effectiveness of EnergyLens for accurate activity inference (appliance, location and the identity of the occupant), we validated it in three different residential scenarios: (a) a controlled residential setting with a single occupant; (b) a controlled shared residence with multiple occupants; and (c) a real residential setting with a single occupant over multiple days. Note, we discuss the challenges of evaluating EnergyLens in a real home setting with multiple occupants in Section 6.

In summary, we make the following contributions:

- We present the design and implementation of **EnergyLens**, a novel algorithm that fuses sensor data from smartphones and energy meter data for activity inference in a residential setting;
- We deployed EnergyLens in both controlled and real world settings, to evaluate different use scenarios. Extensive data collection from these deployments show that EnergyLens is significantly accurate than using the energy meter data alone, particularly in real settings;
- Using empirical data, we extensively discuss the impact of different factors such as simultaneous usage of

appliances with similar power consumption and phone’s orientation, on detection accuracy of EnergyLens and challenges thereof for real world activity classification for energy consumption in buildings.

2. USAGE SCENARIO

We now describe the underlying assumptions and the usage mode of the EnergyLens system. We assume that users have an installed smart meter to measure power usage (real power) for their home. While additional information such as power factor can be further useful to improve EnergyLens accuracy, we do not assume that it is readily available. We further assume that this power data is accessible by the server running EnergyLens over the web.

Training using Smartphones: EnergyLens users are assumed to have a basic Android smartphone with the ability to sample microphone audio and WiFi signal strength. During an initial *training phase*, users run the EnergyLens application and visit each room in their house for a few minutes to provide room-level location annotations. Within each room, the users are required to switch on the appliances they want EnergyLens to identify, wait for some time for its power consumption to reach a steady state and then turn it off. Each of these appliance on-off durations are then annotated, with a recognizable name such as ‘Fan’ or ‘Kettle’ or ‘Microwave’, in the EnergyLens smartphone application. These appliance annotations are then used to train both the audio based classifier (Section 3.1.3) and the meter based NILM algorithm (Section 5). In a home with multiple occupants, where the smartphone models are not identical, this training process has to be repeated for each device *type*. This is due to our observation that microphone audio and WiFi fingerprints vary significantly across different phone models. After the training phase, EnergyLens server learns the necessary models that can then be applied to infer activity during actual usage.

Actual Usage Scenario: As the users start to use EnergyLens in their daily lives, we assume that they mostly carry their phones with them as they move around the house. In other words, EnergyLens assumes that the energy usage events – such as turning an appliance ON or OFF – takes place in the same location as the smartphone’s location. During the activity, microphone audio is sampled and necessary features are extracted from it on the phone itself. Extracted features from the audio stream, along with the data from WiFi scans are sent to the server. Our EnergyLens algorithm executes on the server and identifies *who* performed *what* activity *where* in the home, and *when* they performed it. This annotated energy usage activity can then be presented to the user through the EnergyLens app, or by logging on to a website, or sent to them using a different modality.

3. ENERGYLENS OVERVIEW

At a high level, the goal of EnergyLens is to infer appliance activity and associate it with an occupant by fusing sensory data from smartphones, energy data from smart electricity meters, and metadata information about the appliances in a home. The two main aspects of EnergyLens are therefore the different sensor data sources, and the EnergyLens algorithm that we have developed to perform the activity inferences. We describe these components in detail below.

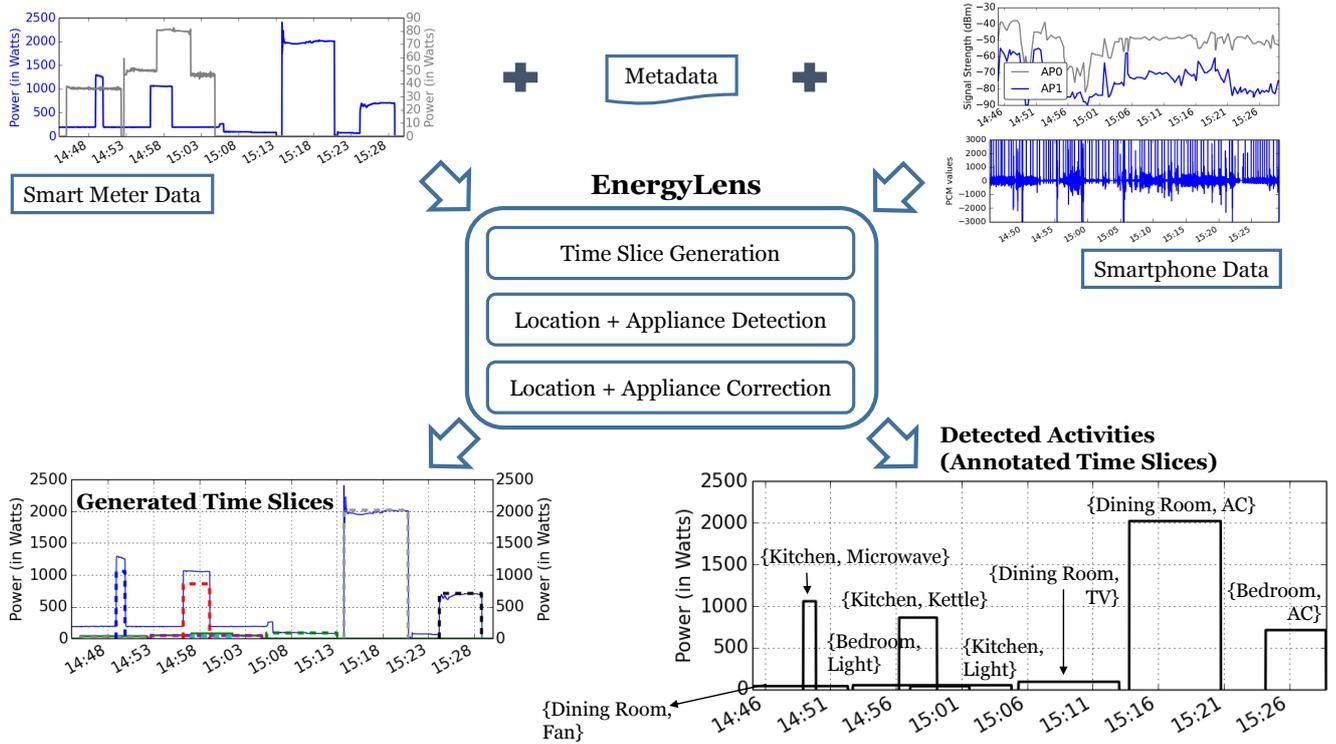


Figure 1: EnergyLens Overview. Smart meter data, appliance metadata and phone data (WiFi scans and audio signals) are taken as inputs. Time Slice Generation (Stage I) generates time slices (shown in dotted lines) from meter data. These time slices are used by Location and Appliance Detection Stages (II and III shown together here). This generates a set of annotated time slices. If location or appliance were misclassified, then Location and Appliance Correction Stages (IV and V) rectify the errors and generate the final set of annotated time slices.

3.1 System Components and Data Sources

3.1.1 Meter data for NILM

We use the total power measurements, sampled at 1Hz, from the whole house electricity meter to create time slices which define a period of activity for each appliance. EnergyLens uses a simple edge detection and edge matching algorithm to retrieve these time slices which are then labeled during the localization and appliance detection stages (see Section 3.2 for details).

3.1.2 WiFi-based localization

WiFi fingerprinting has been a popular technique for indoor localization for many years. EnergyLens uses a simple WiFi fingerprinting technique that has been adapted from RADAR [2]. Our EnergyLens app captures the received signal strengths (RSS) of the nearby access points by doing periodic WiFi scans. For every scan, the phone logs a timestamp and the MAC address, SSID and RSSI values of all the visible APs. These logs are then periodically uploaded to the server for further processing. At the server side, kNN algorithm is used to match and classify the received fingerprints with the stored fingerprints (collected during the training phase) to infer the room-level location for a given time slice.

3.1.3 Audio-based appliance detection

EnergyLens uses the smartphone microphone to capture audio samples when appliances are in use. Collected audio data is then processed by an audio processing pipeline

implemented in our EnergyLens mobile app. This audio processing pipeline has primarily two stages: (1) pre-processing; and (2) feature extraction. The first stage involves *sampling* audio at 8kHz for a period of 10 seconds at every 20 seconds interval (i.e. duty-cycle = 10/20 = 50%). These samples are then combined together into *frames* with a frame size of 500ms. Next, we apply a Hamming window function to each of the frames before passing it on to the next stage. For each of the resulting frames, 13 *Mel-Frequency Cepstral Coefficient (MFCC) features* are computed on the phone. Prior work on acoustic background recognition [19, 28] have shown that MFCC is effective in non-speech audio recognition. These features are then periodically sent to the server for audio classification using the model learned from the training data, while the raw audio data is discarded. Performing the feature computation on the phone helps preserve user privacy, while also significantly reducing the amount of data uploaded to the server.

3.1.4 Metadata

EnergyLens further uses additional information, that is static and can be easily collected. This additional data consists of two different meta information about appliances present in the home:

- Appliance Location Mapping: Room-level location tagged for each appliance.
- Appliance Power Mapping: Measured power for each appliance, extracted from the meter data collected during the training phase (Section 2), is used to tag each

Appliance	Location	Power(W)
AC	Dining Room	1950
Microwave	Kitchen	1150
Kettle	Kitchen	950
AC	Bedroom	670
TV	Dining Room	80
Fan	Dining Room	45
Fan	Bedroom	45
Light	Bedroom	45

Table 1: *Metadata*: Appliance Location Power Mapping

appliance with its power consumption. This information is critical for standard NILM algorithm as well.

Using these two pieces of metadata, we create a table with this information – Appliance Location Power Mapping (hereafter labeled as *Metadata*). A snapshot of the *Metadata* for an example home is shown in Table 1.

3.2 EnergyLens Algorithm

Our proposed EnergyLens algorithm itself comprises of five key stages, as illustrated in Figure 1, with each stage described in detail below.

Stage I: Time Slice Generation

In this stage, the electricity meter data is used to generate time slices i.e. “when” different activities are performed. Each time slice is annotated with its observed power consumption, using meter data, for further stages. The sensor stream for the total real power from the electricity meter is taken as the input. This stage primarily involves two steps - *Edge Detection* and *Edge Matching*.

Edge Detection: In this step, we generate all the rising and falling edges from the meter data stream S_i . Each edge $e_i = (m_i, t_i) \in E_i$ consists of a tuple containing magnitude m_i and time t_i at which the event occurred. m_i is a signed value that captures the power change of an electrical event above a minimum threshold m_t . From the detected edges, we discard edges where:

- A falling edge occurs before a rising edge of similar magnitude.
- Rising and falling edges of similar magnitude occur within a very small time duration. This results in filtering out immediate ON and OFF events.
- Multiple rising (or falling) edges (of similar magnitude) are generated within a time duration. Only the last rising or falling edge is retained amongst such a set while the rest are discarded.

Edge Matching: The filtered edges from the previous step are converted into time slices using a threshold based matching algorithm. We use edge magnitude m_i to match rising and falling edges. Due to an initial surge of power when an appliance is turned on, the magnitude of rising edges are typically higher than the corresponding falling edges. However, in some cases e.g. when there is a heating element that increasingly consumes higher power, the magnitude of falling edges turns out to be higher. Therefore, for each falling edge, the matching algorithm looks for rising edges with magnitude within $\pm p\%$ of the falling edge. Amongst such candidate rising edges for a falling edge, we select the rising edge with the least difference in magnitude with the falling edge. Selected rising and falling edges are

paired to create a time slice $t_s = (t_r, t_f, mag_t)$ where t_r is the start time (time for the rising edge), t_f is the end time (time for the falling edge) and mag_t is the power consumption of the time slice. The final set of time slices T_s , thus created, is passed onto the next stage.

Stage II: Location Detection

In this stage, every time slice in T_s is annotated with a location. For each of the time slices inferred, the WiFi scan data from the phone is used to infer “where” an activity is performed. From the WiFi scan data stream, we first summarize a minute’s worth of data by taking the mean of signal strength samples received from each of the access points (APs) separately. Using these individual summarized values from the different access points, we form a feature vector $\langle rssi_1, rssi_2, \dots, rssi_k \rangle$ where k is the number of visible access points. We ignore APs with less than -85dBm signal strength since they are close to the noise floor. This feature vector is fed into the kNN algorithm to classify every time period (here, a minute) with one of the room locations. Thereafter, a location label which is in majority for the duration of a given time slice is selected as the location for the corresponding time slice. If the duration of the time slice is smaller than the sampling rate of WiFi scan, they are discarded at this stage.

Stage III: Appliance Detection

For all the location annotated time slices, audio recognition is performed to determine the appliance in use. We use an SVM classifier on the 13 MFCC features (extracted from the audio data) received from the phone. Prior work has shown SVM to perform the best for appliance recognition [36, 32]. From the predicted labels for a time slice, the label which is in majority is selected as the appliance being used during that period.

At the end of Stage III, EnergyLens obtains “what” appliance is used, “when” it is used and “where” it is used, which are collectively referred to as an “*activity*”. The identity is implicit for a single occupant setting. For a multiple occupant scenario, “User Association” is done after the Location and Appliance Detection stages (explained in the next subsection). Hereafter, we compare the classified detection for both appliance and location and verify if the combination of classified appliance and location match with the metadata. For all the unmatched entries, we perform the Location and Appliance Correction stages as described below.

Stage IV: Location Correction

In this stage, we use the *Metadata* to extract all the appliances listed to be at the same location as the classified location for a given time slice. For each of these appliances, we compare their power consumption from the *Metadata* with the observed power from the meter. For all time slices where the comparison results do not match (indicator of incorrect location classification), the classified location of the time slice is corrected to the location of the appliance (from amongst all the appliances in the *Metadata*) with the closest match in terms of power consumption. For a non-unique match, if all the matched entries have the same appliance and location (indicating multiple same appliances in the same location), then we use this location for correction. Otherwise, the misclassified location of the time slice is left unchanged.

Stage V: Appliance Correction

We detect an incorrect appliance classification by comparing the power consumption (from *Metadata*) of the classified ap-

Setting	Experiment Type	Appliances Used	Duration
Controlled	Single Occupant	Fans, Lights, AC, TV, Microwave, Kettle	6 hours
Controlled	Multi-Occupant	Fans, Lights, AC, TV, Microwave, Kettle	3 hours
Real	Single Occupant	Lights, TV, Microwave, Kettle	4 days

Table 2: Empirical Dataset Description

Exp#	Description	Appliances Used	Objective
1	Ideal scenario - use of dissimilar appliances with phone kept outside the pocket	Fans, Lights, AC, TV, Microwave, Kettle	Assess algorithm performance
2	Use of dissimilar appliances with phone kept inside the pocket		
3	Use of both similar and dissimilar appliances simultaneously		
4	Realistic scenario - multiple events with combination of similar and dissimilar appliances used simultaneously		
5	Use of similar appliances simultaneously across different rooms	Lights, Fans, TV	Assess impact on time slice accuracy
6	Use of dissimilar appliances used simultaneously across different rooms	Fans, AC, TV, Microwave, Kettle	

Table 3: Controlled experiments for single occupant setting. Here, similar appliances refers to those that lie within the same power consumption range. E.g. Microwave and Kettle lie within 800 – 1200W range; Fans and Lights fall within 30 – 45W range.

pliance with the power consumption observed in the location corrected time slices from the previous stage. For correcting these misclassified appliances, we find all the appliances in the predicted location (using *Metadata*) and select the one with the closest match in terms of power consumption. Again, for a non-unique match, if all the matched entries have the same appliance label then we use this label for correction. Else, we leave the entry with misclassified appliance unchanged.

Multi-Occupant Setting

User Association Stage: For a multi-occupant setting, each time slice further needs to be associated with a user to specify “who” is performing the activity. For each of the generated time slices after Stage I, location and appliance detection (Stages II and III) are performed on data from each of the occupant’s smartphones. Note that each smartphone acts as a proxy for an occupant. At the end of Stage III, each time slice gets multiple sets of <location, appliance> tuples associated with it. The number of occupants determines the number of sets generated for each time slice.

We take the generated set of annotations (each set corresponding to an occupant) for every time slice and match the classified location and inferred power consumption with the *Metadata*. If only one of the sets match with the metadata, then the corresponding occupant is associated with this time slice. If, however, the matching process results in a non-unique match, then we select the set with the least difference from the metadata in terms of power consumption. A non-unique match at the end of this process indicates that the corresponding occupants were all present at the same location where the activity was being done during that period and hence the time slice is associated with all of these occupants.

4. EXPERIMENTAL SETUP

We now describe the experimental setup that we used to evaluate EnergyLens. We conducted experiments in two phases: (1) student volunteers from IIT-Delhi conducting the experiments in controlled settings; (2) involved experiments in a real world setting.

Controlled Experiments

We used a well furnished 3 room apartment with common appliances such as a refrigerator, television, microwave, kettle, room level air conditioning units, multiple lights and fans. All the appliances were used during the experiments. Multiple experiments were performed to emulate the single occupant and the multi-occupant settings. The participants were closely monitored to observe their behavior while performing the activity. They were given a script to follow which mentioned the appliances they had to use and the duration of use. Their actions were not influenced in any way by the authors monitoring these experiments. We collected ground truth manually by annotating each event by the authors as it was being done.

Real World Experiments

The second phase experiments were done in the one of faculty residences at IIT-Delhi. An apartment with a single occupant, consisting of 6 rooms and 8 different types of appliances, was selected. The appliances in the apartment include TV, lights, fans, microwave, room level Air Conditioner (AC), refrigerator, washing machine and kettle. No fans or AC were used during the data collection phase as the experiments were conducted during early winter. The refrigerator was not considered in our approach but it was running in the background for the entire experiment duration of a week. Appliances used during the week were TV, lights, microwave and kettle. The smartphone used by the

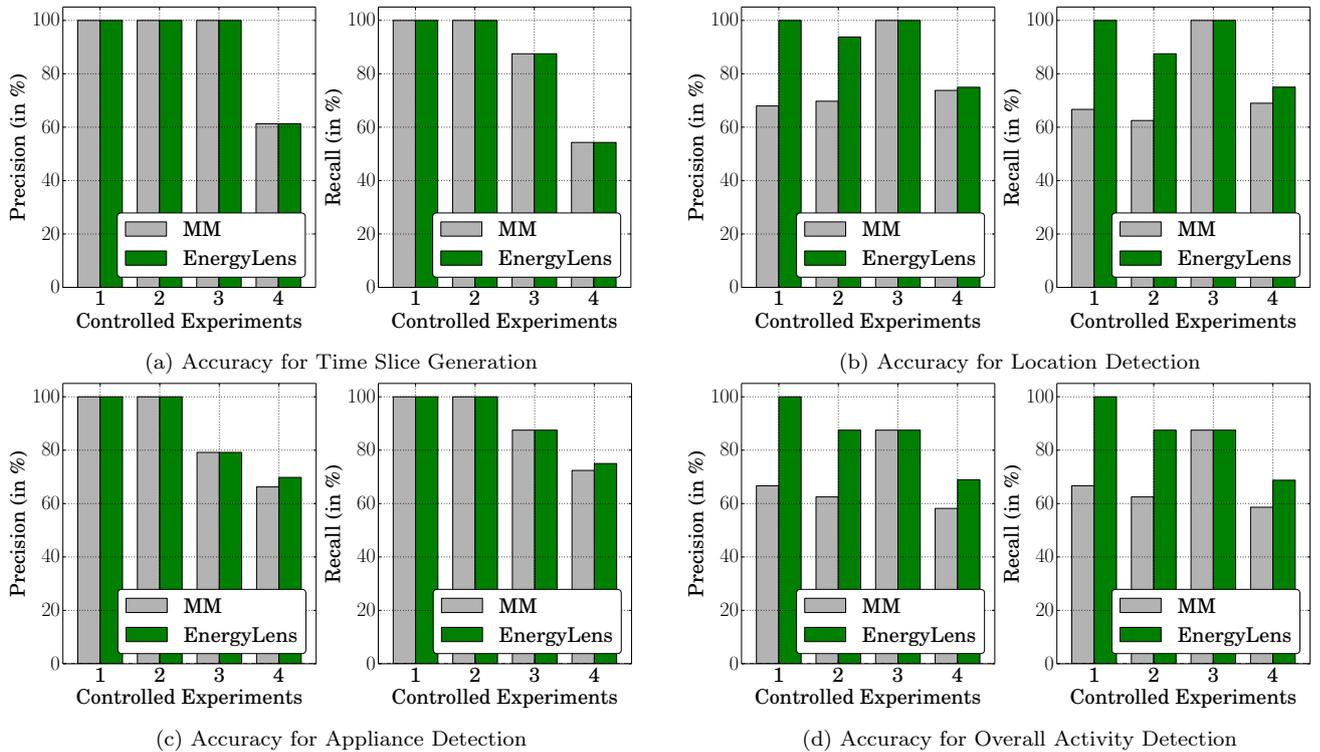


Figure 2: Component wise Precision and Recall for all the controlled experiments (single occupant). Refer Table 3 for experiment details. Here, MM stands for Meter + Metadata (as explained in Section 5).

occupant was a Samsung Galaxy S4 with our EnergyLens application installed. Data description is presented in Table 2. Out of a total 7 experimental days, we obtained 4 days of useful data.

For collecting ground truth, we used the paper-pen method wherein we stuck a paper next to each appliance of interest. The occupant of the home was requested to enter the ON and OFF time as she used appliances during the experiment week. Labeled ground truth was collected every evening from the apartment and was compared with smart meter data for the previous day. Any discrepancies observed in the labels, through manual comparison with smart meter data, were duly discussed with the home occupant and rectified to ensure true validity of the ground truth.

For both phases, the smart meter data was collected at 1Hz. The phone sensors namely, WiFi scan and audio, were sampled every 20 seconds. Audio sampling was performed for 10 seconds every 20 seconds. The phone uploaded data every 5 minutes to the central server. Appliance metadata (appliances, their location and observed power) for both apartments was collected before the experiments and was done by measuring the real power consumption of the appliances as described in Section 2.

5. EVALUATION

We evaluate our EnergyLens system (combining phone sensors with meter data), as described in Section 3.2, by comparing its performance (in terms of activity detection accuracy) with a NILM only approach. For the NILM only approach, which we refer to as **Meter + Metadata (MM)**, *Metadata* information is used in combination with smart

meter power data to determine the appliance and its location. For the MM approach, the time slices are generated (as described in Section 3.2 – Stage I), and are labeled with the appropriate appliance and location by matching the observed power from the meter with the *Metadata*. Activity detection accuracy is measured in terms of:

Precision: calculated as the ratio of the number of activities correctly identified and the total number of identified activities.

Recall: calculated as the ratio of the number of activities correctly identified and the total number of activities done during the experiment duration.

5.1 Performance Analysis

We conducted several experiments, as listed in Table 3, to emulate different realistic scenarios for evaluating the detection accuracy of EnergyLens.

5.1.1 EnergyLens inference evaluation

In this section, we show the activity detection accuracy of EnergyLens together with the accuracy of its individual components, namely, time slice generation (“*when*”), location detection (“*where*”) and appliance detection (“*what*”). We present the results for both controlled and real world experiments in a single occupant setting. To present the component wise accuracy, we use the ground truth after Stage I. In other words, to report the location and appliance detection accuracy, we use ground truth time slices over which we run the detection algorithm and report the accuracy for the individual components.

Day	Approach	Time Slice Generation	Location Detection	Appliance Detection	Activity Detection
1	MM	33.3 / 27.3	66.4 / 62.9	55.3 / 57.9	49.6 / 29.3
	EnergyLens	37.5 / 27.3	90.6 / 87.5	100 / 100	91.7 / 87.5
2	MM	30.7 / 50.0	51.9 / 61.2	62.5 / 50.0	14.3 / 11.2
	EnergyLens	33.3 / 50.0	76.2 / 85.7	100 / 100	76.2 / 85.7
3	MM	70.0 / 77.8	85.0 / 80.0	35.5 / 46.7	22.4 / 26.7
	EnergyLens	70.0 / 77.8	88.1 / 85.7	88.6 / 85.7	68.6 / 71.4
4	MM	70.0 / 63.6	63.7 / 72.0	28.3 / 50.0	27.4 / 22.0
	EnergyLens	77.8 / 63.6	91.1 / 88.9	61.1 / 77.8	64.5 / 66.7

Table 4: Precision/Recall (P/R) of activity detection accuracy for a single occupant real setting. MM stands for Meter + Metadata (as explained in Section 5).

Exp#	Approach	User Association		Activity Detection	
		User 1	User 2	User 1	User 2
1	MM	NA		44.7 / 46.4	75.0 / 75.0
	EnergyLens	80.0 / 80.0	80.0 / 80.0	60.0 / 100	100 / 100
2	MM	NA		100 / 100	25.0 / 25.0
	EnergyLens	66.7 / 50.0	75.0 / 60.0	100 / 100	75.0 / 75.0
3	MM	NA		30.0 / 40.0	66.7 / 66.7
	EnergyLens	83.3 / 83.3	20.0 / 20.0	70.0 / 80.0	80.0 / 80.0

Table 5: Precision/Recall (P/R) of activity detection and user association accuracy for a multi occupant controlled setting. To report activity detection accuracy per occupant for MM approach, the occupant labels are taken from the ground truth.

Figure 2 shows the precision and recall achieved across experiments 1 – 4 (See Table 3). For Experiment 1, wherein the participant complied with all the assumptions made by the algorithm and kept the phone out of their pocket most of the times, EnergyLens achieved 100% precision and recall for activity detection. With phone being kept inside the pocket for Experiment 2, the overall inference accuracy for EnergyLens reduces, primarily due to the poor location detection accuracy caused due to the varying orientation of the phone (as explained in Section 5.1.3). Similarly, in Experiment 3, the accuracy of EnergyLens was lower due to reduced appliance detection accuracy, caused due to mis-classifications by the audio recognition algorithm. Appliance correction stage was unable to rectify these misclassifications as *Metadata* had multiple appliances with similar power consumption at the classified location.

During complex appliance usage events, as emulated in Experiment 4 wherein multiple appliances (combination of similar and dissimilar appliances in terms of power consumption) were used simultaneously across different rooms, EnergyLens is able to achieve 68.93% precision and 68.75% recall which is higher than MM approach which achieves 58.19% precision and 58.23% recall.

Real world experiments for the single occupant setting, as shown in Table 4, further corroborate that EnergyLens always performs at least as good and in most cases much better than the meter only approach (MM). The average precision and recall for EnergyLens was 75.2% and 77.8% while for MM approach, it was 28.4% and 22.3%. MM approach compares the observed power consumption obtained from meter data with the metadata to determine location and appliance. For similar appliances (e.g. lights and fans), which were used extensively by the occupant, MM was not

able to distinguish them. This resulted in low appliance detection accuracy which brought down the activity detection accuracy. In contrast, EnergyLens used phone data for determining the location and appliance which if misclassified was corrected with the help of metadata. This combination of phone with meter improved the activity detection accuracy significantly as seen in Table 4.

5.1.2 User Association Accuracy

In this section, we report the accuracy of the user association component of EnergyLens wherein *Precision* and *Recall* are calculated separately for each user. We only present the results of controlled experiments for the multi-occupant setting. Real world scenarios present a lot of challenges (as discussed in Section 6) which makes the problem of user association non-trivial. Further, obtaining ground truth for such settings is a significant challenge. Detailed instrumentation is required to monitor the occupants’ activities to identify their actual location and appliance activity which we plan to undertake in the future.

In order to test the performance of our user attribution algorithm, we conducted several controlled experiments where we emulated realistic scenarios. Each experiment involved doing a set of activities by two participants based on a script given to them. The results of these experiments are shown in Table 5. The table lists user association accuracy along with the activity detection accuracy for each occupant. For the MM approach, it is evident that meter data cannot be used in isolation to determine the identity of the occupant who did an activity. Therefore, to determine the activity detection accuracy per occupant, we feed pre-annotated user time slices to the MM algorithm. This reduced the multi-occupant case to the single occupant case for this approach.

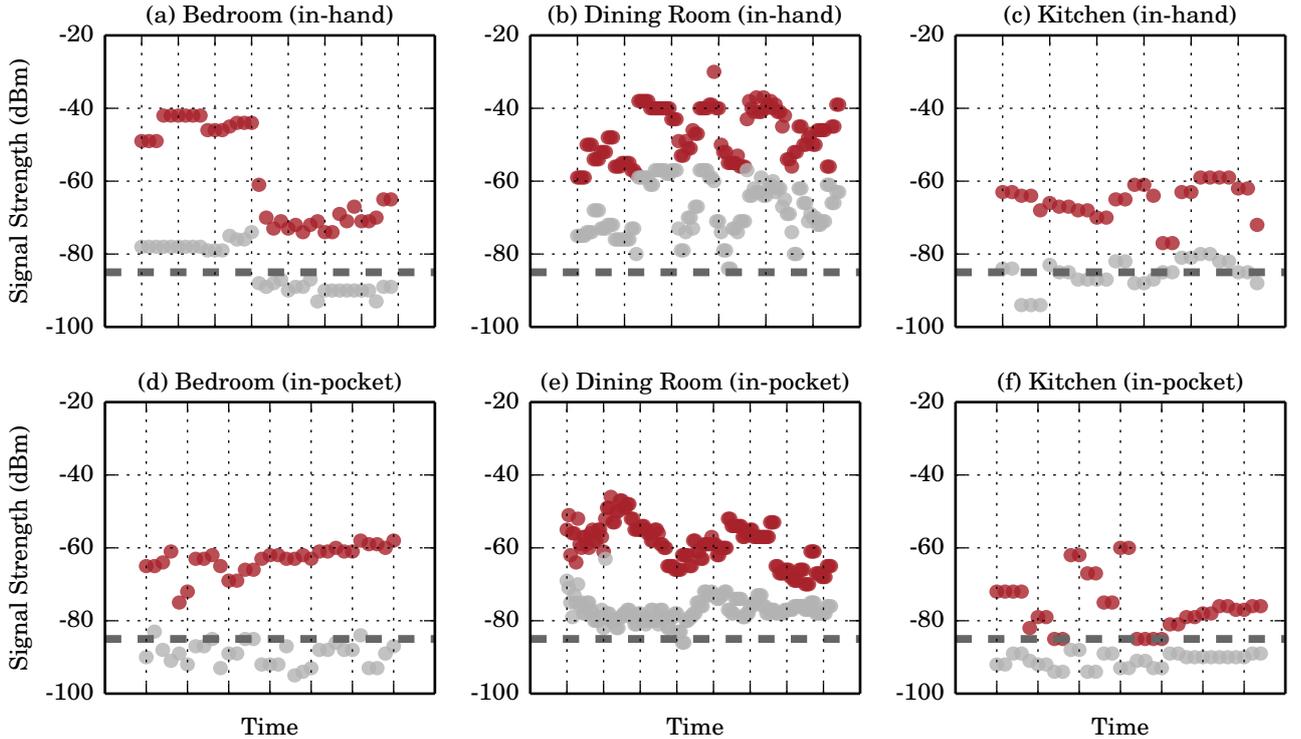


Figure 3: Illustration of signal attenuation caused due to phone orientation. Each color represents the signal strength from a different access point (maroon representing the AP inside the apartment where experiment was conducted). The line drawn at -85dBm denotes that all signal strength values below -85dBm are discarded by EnergyLens.

The results listed in Table 5 look promising and they demonstrate the importance of smartphones in attributing energy to individuals. As pointed out earlier, the accuracy of individual components affect the overall activity detection and user association accuracy. We see the effect of poor location detection accuracy on user association in Experiment 3 for the second occupant. The misclassifications caused by the detection algorithm was due to the poor quality of location data obtained from the phone when it was kept in the pocket during the experiment. The next section (Section 5.1.3) explains the impact of phone’s orientation in detail.

5.1.3 Impact of Phone Orientation on Location Detection Accuracy

In this section, we analyze how the orientation of phone, i.e. phone’s position while performing activities, affects location detection accuracy.

We performed two experiments, referred as Exp 1 and 2 in Table 3. In the first experiment, the phone was carried in hand or placed on the table while doing the activities. In experiment 2, the phone was kept inside the pocket all the time. Four access points (APs) were seen in the vicinity of the apartment; one of them being the AP of the apartment in which experiments were conducted. Out of the four, two APs had less than -85dBm signal strength and were therefore not used by EnergyLens. Two APs were visible across all three rooms. The AP in the apartment was kept in the dining room.

Figure 3 shows the signal strength received from the two visible APs for each room. Observations obtained from experiment 1 and 2 are shown in Figure 3a – Figure 3c and Figure 3d – Figure 3f respectively. Since the apartment AP was in the dining room, the signal attenuation is negligible and thus doesn’t affect the detection accuracy for this room. However, signals from the second AP which were comparatively weaker in the Bedroom and Kitchen, gets further attenuated due to phone’s orientation and affects the localization accuracy for these rooms.

The experiments were conducted in an apartment building inside IIT-Delhi campus in India where very few access points were seen in the vicinity. The accuracy is expected to be much more when the number of access points increase in number (which is relatively common in the US).

5.1.4 Impact of Sampling Intervals on Appliance Detection Accuracy

To assess the impact of the sampling interval on the accuracy of appliance detection, we conducted an experiment where multiple appliances were used (experiment 4 in Table 3) with audio sampling done every 20 seconds for a period of 10 seconds. We sub-sampled this data to obtain data with sampling intervals from 40 seconds to 180 seconds. Table 6 shows the appliance detection accuracy of the algorithm for different sampling intervals. Since the usage of appliances during this experiment was well separated in terms of power consumption and usage times, and 10 second audio data was observed to be sufficient for classification

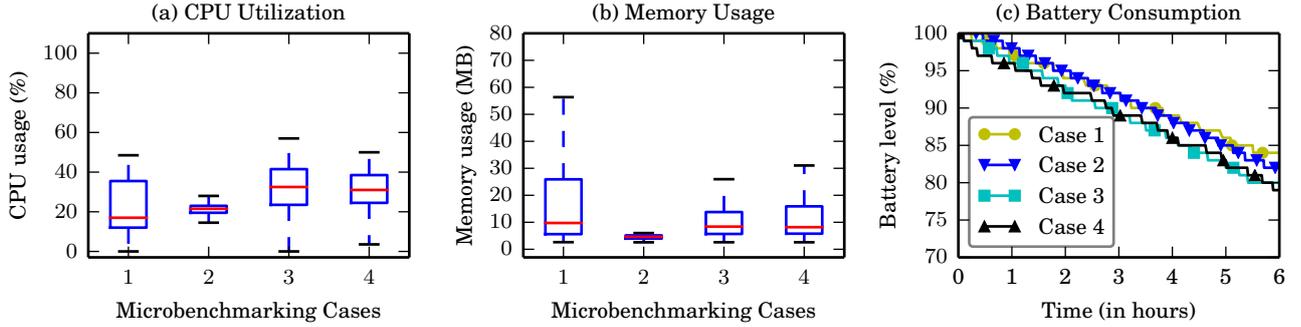


Figure 4: Microbenchmarking of EnergyLens in terms of CPU, memory and battery usage. The four cases are: (1) Wifi sampling only (2) Audio Sampling only (3) Wifi and Audio sampling together (4) EnergyLens with sampling and data upload

Sampling Interval(s)	EnergyLens	
	Precision(%)	Recall(%)
20	69.8	75.0
40	66.0	68.8
60	67.3	71.9
80	67.2	71.9
100	68.9	75.0
120	64.1	65.6
140	59.3	62.5
160	62.8	62.5
180	64.1	65.6

Table 6: Impact of Sampling Intervals on Precision/Recall of Appliance Detection Accuracy

purposes, the appliance detection accuracy remains approximately constant across different sampling intervals. This shows that for such infrequent appliance usages (which are common in real world settings as well), we can reduce the sampling interval for audio sampling without significantly affecting the inference accuracy.

5.1.5 Impact of Simultaneous Activity on Time Slice Generation Accuracy

Simple edge detection and edge matching algorithm for inferring usage durations, as used in the current implementation of EnergyLens, are accurate when the used appliances differ in terms of their power consumption. However, when appliances with similar power consumption (such as fans and lights; microwave and electric kettle) are used simultaneously then it becomes difficult for the edge matching algorithm to correctly associate an ON event with a corresponding OFF event.

We conducted three controlled experiments (refer Table 3): first with dissimilar appliances (fan, microwave, TV, AC, kettle) working at the same time (experiment 6), second with multiple similar appliances (mostly lights and fans) used simultaneously across rooms (experiment 5), and the third with a combination of similar and dissimilar appliances used concurrently (experiment 3). These experiments yielded 100/80, 62.5/45.5 and 100/87.5 of precision/recall respectively, illustrating that the current time slice generation algorithm finds it difficult to differentiate between appliances

with similar power consumption. We believe that additional use of power factor will help differentiate between appliances of similar power consumption and will significantly improve the accuracy for such usage scenarios.

5.2 EnergyLens Application Profiling

In order to characterize the system and energy impact of EnergyLens mobile application, we conducted tests to benchmark the impact of different components of EnergyLens. We used a Sony Xperia SP smartphone with 1.7 GHz dual core processor and 1 GB RAM for the evaluation. The audio sampling rate was set to 8kHz and was sampled for 10 seconds at an interval of 20 seconds. The phone also scanned for WiFi APs every 20 seconds. The collected and processed data was uploaded every 5 minutes to the central server.

5.2.1 Application Overhead

We collected system’s total CPU and memory usage for a duration of 3 hours each, for the cases: (a) System with Wifi turned ON and without EnergyLens installed (baseline case) (b) System with EnergyLens installed and background data collection process enabled. Background running processes were constant for both cases. Mean memory usage of the system for the baseline case was found to be 240MB and 256MB for EnergyLens. The memory overhead of EnergyLens was observed to be 16MB, only a 6% increase in usage. The mean CPU usage was 13% and 22% for cases (a) and (b) respectively. Therefore, the difference in usage (CPU overhead of EnergyLens) was found to be 9%. This increase in CPU utilization of EnergyLens can be reduced with adaptive sampling techniques which we leave as a future work.

5.2.2 Microbenchmarking

The microbenchmarking tests were conducted to quantify the effect of the individual components of EnergyLens, namely Wifi sampling, audio sampling and data upload. We measured the CPU, memory and battery usage for four cases: (1) App with Wifi sampling only (2) App with audio sampling only (3) App with Wifi and audio sampling together (4) EnergyLens App with audio and Wifi sampling with periodic data uploads. CPU and memory measurements were taken every second using Android’s *adb* utility over a period of 3 hours each. Battery runs were done over a period of 6 hours each where the residual battery level

was logged once every 20 seconds by the mobile application during data sampling.

CPU and Memory Footprint: Figure 4a and Figure 4b show the CPU and memory usage of the smartphone for each of the four cases. The mean memory usage is 17MB, 4MB, 10MB and 11MB for cases 1 – 4 respectively and the mean CPU usage of the four cases is 23%, 21%, 32% and 30% respectively. From Figure 4a and Figure 4b, we see that the audio sampling (case 2) component’s impact is less than Wifi sampling component (case 1) in terms of CPU and memory usage both. The upload component (case 4) doesn’t add much memory overhead (only 1MB) when compared to case 3. Similarly, the mean CPU usage of case 3 and case 4 also remains the same. Thus, we can conclude that the additional periodic uploads (case 4) doesn’t add any significant overhead on the system’s performance when compared to case 3.

Energy footprint: The battery drain was observed to be 2.67%, 3%, 3.33% and 3.5% per hour for cases 1 – 4 respectively. They correspond to 37.4 hours, 33.3 hours, 30.0 hours and 28.5 hours of battery life. From Figure 4c, we observe that the difference in battery usage for case 3 and case 4 is not remarkable (just 5% increase from case 3) which reaffirms that the upload component doesn’t add any significant overhead on the system’s performance. On the contrary, the audio sampling (case 2) component’s battery usage is 12% more than Wifi sampling component (case 1), making it the most power hungry component of EnergyLens (case 4). With the battery lifetime of approximately 28.5 hours for the smartphone with EnergyLens application running in the background, we can conclude that EnergyLens is comparable to the normal usage scenarios and it does not reduce battery life of a device perceptibly.

6. DISCUSSION AND LIMITATIONS

EnergyLens seeks to address a very complex problem of identifying and disaggregating energy usage per appliance and associating each usage with a particular user. The current implementation of EnergyLens tries to achieve this objective with a basic set of sensors from mobile phones and home level electricity meters, chosen to keep the overall cost of deployment and maintenance low. We now discuss several challenges that arise while using such a restricted set of sensing modalities for this complex problem.

We note that energy attribution in a multi-occupant setting can not be done by using meter data alone and therefore adding phone information becomes critical. We further show that, by combining even the most simplistic algorithms for localization, appliance detection and NILM, EnergyLens can provide better accuracy for activity inference as compared to NILM only techniques. A natural extension of our work is to investigate more complex algorithms, such as Factorial HMM for NILM [24], sound separation techniques such as Computational Auditory Scene Analysis (CASA) [3] for audio-based appliance detection and localization techniques such as UnLoc [31] and Ariel [12], to further improve accuracy, particularly for more complex scenarios such as in multi-occupant homes. Use of additional sensors for both smartphones and energy meters can provide additional contextual information that can be used to further improve the accuracy. Examples of these additional context information include detection of ambient light using the phone light sensor, improved location classification using camera [1], radio

ranging such as the one provided by Bluetooth Low Energy and power factor to help distinguish loads. Finally, a simple extension of EnergyLens can be to use events on the meter data to trigger sensing for mobile phone audio and WiFi, thereby reducing the overall energy burden due to EnergyLens.

In terms of deployment and use of our system, there are a number of practical considerations that affect the overall inference accuracy. Examples of these considerations include phone orientation affecting WiFi-based localization accuracy, scenarios when the user is not carrying the phone with them while performing any appliance usage and scenarios wherein the user initiates an appliance activity (e.g. start a microwave) and then leaves the room to return back after a time when the activity is still underway. Such considerations are further complicated for the multi-occupant scenario. Examples of complex multi-occupant scenarios include simultaneous switching of appliances by different users across different rooms, movement of users across different rooms while the appliance usage (e.g. television being ON) in each of the rooms is ongoing and one user’s phone is in a room where the second user is present, though without her phone, while initiating an appliance usage activity. These complex scenarios, for both single and multi-occupant case, potentially require both additional sensing information and more sophisticated inference algorithms that also decide when to use or ignore different context information.

Since EnergyLens uses smartphone sensors such as audio from the microphone, privacy is a natural concern. Pre-processing the audio on the phone itself while sending only the inferred values to the EnergyLens server for analysis, mitigates such privacy concerns in EnergyLens. Furthermore, the data analysis for activity inference is simple enough to be easily run on a local computer in the occupants’ home, further alleviating any privacy concerns.

7. RELATED WORK

Location based services using WiFi on phones as a sensor (e.g. LifeMap [4] and ParkSense [20]) as well as acoustic sensing applications using phones (e.g. SpeakerSense [16] and SocioPhone [15]) have been proposed. Smartphones in a way provide an ideal platform to host many of these applications. Motivated by the ubiquity of smartphones and their increasing sensing capabilities, EnergyLens looks at a previously unexplored dimension of residential energy monitoring by leveraging smartphone sensors together with smart meter data for improved electrical appliance activity detection accuracy. We now present related work across the three dimensions of EnergyLens: location inference, audio classification and appliance detection. To the best of our knowledge, there does not exist any other work that addresses the complex combination of inference including “who”, “what”, “when” and “where” for activities in residential setting.

Several techniques, other than the WiFi fingerprinting methods, have been proposed over the years for indoor localization. These techniques use additional sensors e.g. microphone sound, camera [1] and light [25] to improve the overall localization accuracy. EnergyLens can potentially improve its localization accuracy by utilizing such additional smartphone sensors and incorporating some of these sophisticated techniques rather than simple nearest neighbor approach based on WiFi fingerprinting, currently used in this work.

Several frameworks have been proposed for building context aware applications using audio as a sensing modality. These include SoundSense [17] for distinguishing voice and music from ambient noise, Jigsaw [18] for human activity inferences and iSleep [9] for detecting sleep quality. Recently, Auditeur [21], a general purpose acoustic event detection platform for smartphones was proposed that enable applications to register for receiving notifications for the desired audio events as classified within their framework. Several of the techniques used in these frameworks for audio processing on the phone can also be incorporated into EnergyLens for improving the overall detection accuracy.

Energy monitoring in a residential setting involves multiple facets e.g. NILM work using smart meters for energy monitoring and disaggregation [10, 7, 14]; use of additional sensing modalities that indirectly monitors the environment for effects caused due to energy usage [8]; combining meters with other ambient sensors such as light, temperature, motion and acoustic signals [13, 23]; and very recently using smartphones only [30, 29]. EnergyLens does not require deploying and managing additional sensors but endeavors to better solve the problem of energy monitoring by leveraging smartphone sensors for improved appliance detection and recognition.

Of the recent work, AppliSense [33] has some aspects that are related to our work. AppliSense, propose a NILM algorithm to disaggregate appliance energy use. They also propose using a smartphone UI to help users label appliances as part of their training phase. The crucial differentiator is that they don't actually propose, or use, any of the smartphone sensors for the actual disaggregation algorithm. In contrast, the key contribution of EnergyLens is the idea of combining sensor readings from smart meters and user's smartphones for improved energy use inference in order to answer the *who, what, where* and *when* questions of energy apportionment. Kay et al. [11] presented the case for apportionment in a building. They studied different apportionment policies and how sensor systems could be used to further enable this. EnergyLens takes a step further and proposes an actual system to provide the key pieces of information required for apportioning energy in homes.

8. CONCLUSIONS

This paper presents the design and implementation of EnergyLens, a system for accurate activity inference in homes – specifically to disaggregate energy and apportion usage to the occupants. The key contribution of EnergyLens is a novel sensor fusion algorithm that utilizes sensor data from occupants' smartphones and that from a smart meter to improve inference accuracy as compared to traditional NILM approaches. In contrast to prior approaches that propose to address this problem using additional sensors and complex algorithms, we show that even simple NILM, audio-based classification and WiFi-based localization methods provide sufficient accuracy in realistic settings.

We highlight several challenges that impact the inference accuracy for the complex problem at hand. These include the impact of phone orientation, diverse usage scenarios especially for the multi-occupant case, collection of ground truth data for establishing the accuracy and simultaneous switching of appliances with similar power consumption profiles. While the current implementation of EnergyLens seeks reasonable accuracy in many such complex scenarios, we

plan to enrich our system with additional sensing modalities both from the smartphone (e.g. light sensor, camera) and the smart meter (e.g. power factor) together with more sophisticated inference algorithms. The flip side, however, of using sophisticated inference algorithms will be the increased computation demand; EnergyLens is currently light-weight enough to run on a laptop class machine.

9. ACKNOWLEDGEMENTS

Authors would like to acknowledge the support provided by ITRA project, funded by DEITY, Government of India, under grant with Ref. No. ITRA/15(57)/Mobile/HumanSense/01. This work was also supported in part by NSF grant SHF-1018632. We wish to thank our anonymous reviewers for their feedback to help improve the paper.

10. REFERENCES

- [1] M. Azizyan, I. Constandache, and R. Roy Choudhury. SurroundSense: Mobile phone localization via ambient fingerprinting. In *Proc. of the 15th Annual International Conference on Mobile computing and networking*, pages 261–272. ACM, 2009.
- [2] P. Bahl and V. N. Padmanabhan. RADAR: An in-building RF-based user location and tracking system. In *Proc. of the 19th Annual Joint Conference of the IEEE Computer and Communications Societies*, volume 2, pages 775–784. IEEE, 2000.
- [3] G. J. Brown and M. Cooke. Computational auditory scene analysis. *Computer Speech & Language*, 8(4):297–336, 1994.
- [4] J. Chon and H. Cha. Lifemap: A smartphone-based context provider for location-based services. *IEEE Pervasive Computing*, 10(2):58–67, 2011.
- [5] DOE. Buildings Energy Data Book. *Department of Energy*, March 2009.
- [6] J. Froehlich, L. Findlater, and J. Landay. The design of eco-feedback technology. In *Proc. of the CHI Conference on Human Factors in Computing Systems*, pages 1999–2008. ACM, 2010.
- [7] H. Gonçalves, A. Ocneanu, M. Bergés, and R. Fan. Unsupervised disaggregation of appliances using aggregated consumption data. In *The 1st KDD Workshop on Data Mining Applications in Sustainability*, 2011.
- [8] S. Gupta, M. S. Reynolds, and S. N. Patel. ElectriSense: Single-point Sensing using EMI for Electrical Event Detection and Classification in the Home. In *Proceedings of the 12th ACM international conference on Ubiquitous Computing (Ubicomp)*, pages 139–148. ACM, 2010.
- [9] T. Hao, G. Xing, and G. Zhou. iSleep: unobtrusive sleep quality monitoring using smartphones. In *Proc. of the 11th ACM Conference on Embedded Networked Sensor Systems*, page 4. ACM, 2013.
- [10] G. W. Hart. Nonintrusive appliance load monitoring. *Proceedings of the IEEE*, 80(12):1870–1891, 1992.
- [11] S. Hay and A. Rice. The case for apportionment. In *Proc. of the 1st ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Buildings*, pages 13–18. ACM, 2009.
- [12] Y. Jiang, X. Pan, K. Li, Q. Lv, R. P. Dick, M. Hannigan, and L. Shang. Ariel: Automatic wi-fi

- based room fingerprinting for indoor localization. In *Proc. of the 14th ACM Conference on Ubiquitous Computing*, pages 441–450. ACM, 2012.
- [13] Y. Kim, T. Schmid, Z. Charbiwala, and M. Srivastava. ViridiScope: design and implementation of a fine grained power monitoring system for homes. In *Proc. of the 11th International Conference on Ubiquitous computing*, pages 245–254. ACM, 2009.
- [14] J. Z. Kolter, S. Batra, and A. Ng. Energy Disaggregation via Discriminative Sparse Coding. In *Advances in Neural Information Processing Systems*, pages 1153–1161, 2010.
- [15] Y. Lee, C. Min, C. Hwang, J. Lee, I. Hwang, Y. Ju, C. Yoo, M. Moon, U. Lee, and J. Song. SocioPhone: Everyday Face-to-Face Interaction Monitoring Platform Using Multi-Phone Sensor Fusion. In *Proc. of the 11th International Conference on Mobile Systems, Applications, and Services*. ACM, 2013.
- [16] H. Lu, A. B. Brush, B. Priyantha, A. K. Karlson, and J. Liu. SpeakerSense: energy efficient unobtrusive speaker identification on mobile phones. In *Pervasive Computing*, pages 188–205. Springer, 2011.
- [17] H. Lu, W. Pan, N. D. Lane, T. Choudhury, and A. T. Campbell. SoundSense: scalable sound sensing for people-centric applications on mobile phones. In *Proc. of the 7th International Conference on Mobile systems, applications, and services*, pages 165–178. ACM, 2009.
- [18] H. Lu, J. Yang, Z. Liu, N. D. Lane, T. Choudhury, and A. T. Campbell. The Jigsaw continuous sensing engine for mobile phone applications. In *Proc. of the 8th ACM Conference on Embedded Networked Sensor Systems*, pages 71–84. ACM, 2010.
- [19] L. Ma, B. Milner, and D. Smith. Acoustic environment classification. *ACM Transactions on Speech and Language Processing*, 3(2):1–22, 2006.
- [20] S. Nawaz, C. Efstratiou, and C. Mascolo. ParkSense: a smartphone based sensing system for on-street parking. In *Proc. of the 19th Annual International Conference on Mobile Computing & Networking*.
- [21] S. Nirjon, R. F. Dickerson, P. Asare, Q. Li, D. Hong, J. A. Stankovic, P. Hu, G. Shen, and X. Jiang. Auditeur: A mobile-cloud service platform for acoustic event detection on smartphones. In *Proc. of the 11th International Conference on Mobile systems, applications, and services*. ACM, 2013.
- [22] S. N. Patel, T. Robertson, J. A. Kientz, M. S. Reynolds, and G. D. Abowd. At the flick of a switch: Detecting and classifying unique electrical events on the residential power line. In *Proc. of the 9th International Conference on Ubiquitous computing*, pages 271–288. Springer, 2007.
- [23] D. E. Phillips, R. Tan, M.-M. Moazzami, G. Xing, J. Chen, and D. K. Yau. Supero: A sensor system for unsupervised residential power usage monitoring. In *IEEE International Conference on Pervasive Computing and Communications*, volume 18, pages 66–75, 2013.
- [24] D. Rahayu, B. Narayanaswamy, S. Krishnaswamy, C. Labbe, and D. P. Seetharam. Learning to be energy-wise: Discriminative methods for load disaggregation. In *3rd International Conference on Future Energy Systems*, pages 1–4, 2012.
- [25] N. Ravi and L. Iftode. Fiatlux: Fingerprinting rooms using light intensity. In *Adjunct Proc. of the Fifth International Conference on Pervasive Computing*. Citeseer, 2007.
- [26] A. Rowe, M. Berges, and R. Rajkumar. Contactless Sensing of Appliance State Transitions Through Variations in Electromagnetic Fields. In *Proc. of the 2nd ACM Workshop on Embedded Systems For Energy-Efficient Buildings*, 2010.
- [27] V. Srinivasan, J. Stankovic, and K. Whitehouse. FixtureFinder: discovering the existence of electrical and water fixtures. In *Proceedings of the 12th International Conference on Information Processing in Sensor Networks*, pages 115–128. ACM, 2013.
- [28] Z. C. Taysi, M. A. Guvensan, and T. Melodia. TinyEars: spying on house appliances with audio sensor nodes. In *Proc. of the 2nd ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Buildings*, pages 31–36. ACM, 2010.
- [29] M. Uddin and T. Nadeem. EnergySniffer: Home energy monitoring system using smart phones. In *Proc. of 8th International Wireless Communications and Mobile Computing Conference*, pages 159–164. IEEE, 2012.
- [30] M. Uddin and T. Nadeem. MachineSense: detecting and monitoring active machines using smart phone. *Mobile Computing and Communications Review*, 16(4):16–17, 2012.
- [31] H. Wang, S. Sen, A. Elgohary, M. Farid, M. Youssef, and R. R. Choudhury. No need to war-drive: Unsupervised indoor localization. In *Proc. of the 10th International Conference on Mobile systems, applications, and services*, pages 197–210. ACM, 2012.
- [32] J.-C. Wang, H.-P. Lee, J.-F. Wang, and C.-B. Lin. Robust environmental sound recognition for home automation. *IEEE Transactions on Automation Science and Engineering*, 5(1):25–31, 2008.
- [33] M. Weiss, A. Helfenstein, F. Mattern, and T. Staake. Leveraging smart meter data to recognize home appliances. In *IEEE International Conference on Pervasive Computing and Communications*, pages 190–197. IEEE, 2012.
- [34] T. Weng, B. Balaji, S. Dutta, R. Gupta, and Y. Agarwal. Managing Plug-Loads for Demand Response within Buildings. In *Proc. of the 3rd ACM Workshop on Embedded Sensing Systems For Energy-Efficiency In Buildings*, 2011.
- [35] J.-P. Zimmermann, M. Evans, J. Griggs, N. King, L. Harding, P. Roberts, and C. Evans. Household Electricity Survey. A study of domestic electrical product usage. Technical Report R66141, DEFRA, May 2012.
- [36] A. Zoha, A. Gluhak, M. Nati, M. A. Imran, and S. Rajasegarar. Acoustic and device feature fusion for load recognition. In *Proc. of the 6th IEEE International Conference Intelligent Systems*, pages 386–392. IEEE, 2012.