

Non-Invasive Sensor Solutions for Activity Recognition in Smart Homes

A Dissertation

Presented to

the Faculty of the School of Engineering and Applied Science

University of Virginia

In Partial Fulfillment

of the requirements for the Degree

Doctor of Philosophy (Computer Science)

by

Vijay Srinivasan

May 2012

Abstract

Smart home sensor systems that infer residents' activities enable a number of exciting medical monitoring and energy conservation applications. Existing home activity recognition systems are *invasive*, since they require significant manual effort from the end user in installing or training the system, inconvenience the resident by requiring them to constantly wear tags, or require invasive cameras or expensive sensors. Our main hypothesis is that by effectively using data fusion techniques, leveraging the existing smart meter infrastructure in homes, and using only weak biometric sensing, we can build convenient, accurate home activity recognition solutions for the end user.

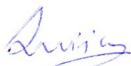
The key challenges for home activity recognition systems addressed in this dissertation include reducing configuration effort from the end user, reducing sensor installation effort, and identifying residents in multi-person homes without using invasive sensors. To reduce user configuration effort, we develop an unsupervised activity recognition algorithm called AUTOLABEL that leverages data fusion and cross-home activity models to accurately recognize resident activities without user training. To eliminate many direct sensors in the home for activity recognition, we develop effective Bayesian data fusion techniques, which combine the existing smart meter infrastructure in homes, with one low cost, non-invasive sensor per room. In our WaterSense approach, we combine a single smart water meter per home, with an occupancy sensor per room to eliminate direct sensors on individual water fixtures. In our LightSense approach, we combine a single smart energy meter per home with a light sensor per room, to eliminate direct sensors on individual light fixtures. Finally, we propose the use of resident height, which is a weak biometric, in an effective data fusion approach, to identify residents for activity recognition in multi-person homes.

We evaluate our proposed activity recognition solutions through short term prototype sensor deployments in home environments lasting from 7 to 10 days each. We show that our low cost,

convenient solutions satisfy the activity recognition needs of numerous smart home applications, such as remote medical monitoring for elderly residents, and fine-grained resource consumption monitoring of light and water fixtures in the home. Finally, we observe that our unsupervised activity recognition algorithm can be used in a wireless snoop attack on smart homes, to infer the residents' daily activities with high accuracy in spite of encrypted wireless transmissions. We propose and evaluate a suite of privacy solutions to mitigate the inference accuracy of such an attack without affecting the performance or functionality of the home activity recognition system.

Approval Sheet

This dissertation is submitted in partial fulfillment of the requirements for the degree of
Doctor of Philosophy (Computer Science)



Vijay Srinivasan

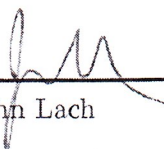
This dissertation has been read and approved by the Examining Committee:



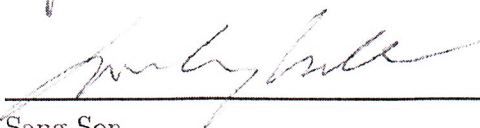
John A. Stankovic, Advisor



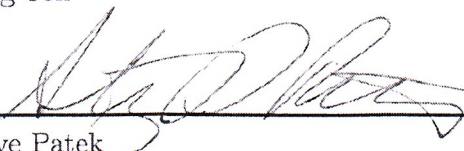
Kamin Whitehouse, Committee Chair



John Lach

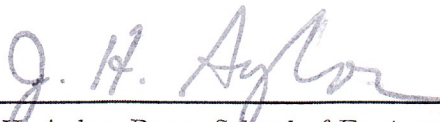


Sang Son



Steve Patek

Accepted for the School of Engineering and Applied Science:



James H. Aylor, Dean, School of Engineering and Applied Science

May 2012

Contents

Contents	iv
List of Tables	vii
List of Figures	viii
1 Introduction	1
1.1 Background	2
1.2 Outline of our activity recognition approach	3
1.3 Privacy considerations	5
1.4 Limitations	6
1.5 Summary of Contributions	8
2 Background	11
2.1 Sensing Approaches	12
2.1.1 Video and Audio based sensing	12
2.1.2 Invisible man approach to activity sensing	12
2.2 Inference approaches	14
2.2.1 Supervised learning	15
2.2.2 Unsupervised learning	15
3 Reducing user configuration effort	17
3.1 AUTOLABEL design	18
3.1.1 Tier-0: General Activity Detection	19
3.1.2 Tier-I Clustering	20
3.1.3 Tier-II Room Classification	22
3.1.4 Tier-III Sensor and Activity Classification	23
3.2 Experimental setup	25
3.2.1 Home Sensor Deployments	25
3.2.2 Evaluation Metrics for Activity Recognition	26
3.3 Evaluation	27
3.3.1 Tier-I Evaluation	27
3.3.2 Tier-II Evaluation	28
3.3.3 Tier-III Evaluation	29
3.4 Limitations	30
3.5 Conclusions	32

4	Reducing sensor installation effort to monitor water fixture usage	34
4.1	State of the art	37
4.2	WaterSense System Design	38
4.2.1	Physical Sensor Components	39
4.2.2	Tier I: Detecting Water Flow Events	40
4.2.3	Tier II: Creating Room Clusters	41
4.2.4	Tier III: Differentiating Fixture Types	44
4.3	Evaluation	44
4.3.1	Experimental setup	44
4.3.2	Fixture monitoring accuracy	45
4.4	Limitations	49
4.5	Conclusions	50
5	Reducing sensor installation effort to monitor electrical fixture usage	51
5.1	State of the art	55
5.2	LightSense Design	57
5.2.1	Tier I: Light and Power Edge detection	57
5.2.2	Tier II: Data Fusion and Matching	62
5.2.3	Tier III: Fixture Identification	67
5.3	Experimental setup	67
5.3.1	Real World Deployments	68
5.3.2	Energy Usage Ground Truth	68
5.4	Assessing Energy Usage	69
5.5	Analysis	72
5.5.1	Impact of Data Fusion and Matching	72
5.5.2	Impact of Bayesian Clustering	74
5.5.3	Parameter Sensitivity Analysis	75
5.6	Limitations	76
5.7	Conclusions	77
6	Non-Invasive Resident Identification	79
6.1	State of the art	82
6.2	Overview: Sensing Height with Ultrasonic Sensors	85
6.3	Experiments in a Controlled Lab Environment	86
6.3.1	Experimental setup	86
6.3.2	Evaluation results	86
6.4	Experiments in Home Environments	89
6.5	Accuracy of Height Sensing in Homes Nationwide	92
6.6	Improving Height Measurement Accuracy with History	93
6.7	Limitations	96
6.8	Conclusions	98
7	Wireless privacy for home activity recognition systems	100
7.1	Background	102
7.2	The FATS Attack	104
7.3	Design Guidelines to Enhance Privacy	106
7.3.1	Using Signal Attenuators	106
7.3.2	Using Random Delays	108

7.3.3	Using Periodic Transmission	109
7.3.4	Using Fingerprint Masking	111
7.3.5	Introducing Spurious or Fake Transmissions	112
7.3.6	Hybrid Schemes	113
7.4	Limitations	114
7.5	Conclusions	114
8	Conclusions	116
8.1	Key contributions towards home activity recognition systems	117
8.1.1	Lower training and configuration effort	117
8.1.2	Fewer, easier to install, low cost sensors	117
8.1.3	More convenient resident identification	118
8.1.4	Improved wireless privacy	119
8.2	Example applications enabled by our activity recognition approach	119
8.3	Contributions beyond smart homes	120
8.4	Future improvements	121
	Bibliography	123

List of Tables

3.1	Characteristics of the eight homes used in our AUTOLABEL deployments	26
6.1	Details of the 3 homes used in deployments to evaluate height based resident identification	91

List of Figures

3.1	The AUTOLABEL Activity Recognition Algorithm that uses just the timestamp T , and unique source identifier ID, of each binary sensor firing to automatically infer daily activities of residents without any configuration or training effort.	19
3.2	Sample data from one of our home sensor deployments. The textual labels and logical groupings are derived from the raw sensor data (black lines) by our AUTOLABEL algorithm.	20
3.3	Examples of the sensors we deployed to evaluate AUTOLABEL.	26
3.4	Accuracy of clustering sensors into rooms in AUTOLABEL, using both <i>dbscan</i> and <i>k</i> -means with 4 different values of <i>k</i>	28
3.5	Overall Inference Accuracy of AUTOLABEL on events detected in Tier 0, Tier II, and Tier III, using all three metrics: EDR, TPR, and DA.	29
4.1	Three tier inference algorithm used by WaterSense to find 1) water flow events 2) clusters of flow volumes that often co-occur with certain motion sensors, and 3) different types of fixtures in the same room.	37
4.2	Sensors used in WaterSense and its evaluation. WaterSense uses a single flow sensor at the water mains (a) and motion sensors (b). Sensors were placed on fixtures for evaluation purposes only (c).	39
4.3	A 6 minute data trace from Home 1 of the WaterSense deployments that shows the water flow data and binary occupancy data from two simultaneous flush events. Tier I detects four edges with flow rates of 0.3 and 0.6 kl/hour, respectively, and identifies them as two different flow events. Tier II assigns the two water flow events to two different bathrooms based on their temporal proximity to motion sensor data in the two rooms, as well as the historical correlation between these flow volumes and motion sensors. Tier III identifies both flow events as toilets using the flow volume and duration.	40
4.4	ELoc Bayesnet used in WaterSense to cluster water flow events based on both flow volume and co-occurrence with motion sensors.	43
4.5	Home deployment details for WaterSense. Our deployments involved two homes with multiple residents. Both homes had a kitchen sink and two bathrooms with both sink and toilet.	45
4.6	Confusion matrix of WaterSense when classifying water flow events across the two homes. Confusion between fixtures of the same type in different rooms is common due to overlapping occupancy. Confusion between fixtures of different types, such as a sink and a flush, due to overlapping flow signatures, is less common. (B stands for bathroom, K for kitchen, S for sink, and F for flush)	46

4.7	Water consumption feedback provided by WaterSense for the various fixtures monitored in the two homes, along with closely matching ground truth consumption. End users can use such a display to start considering cost effective ways to conserve water. (B stands for bathroom, K for kitchen, S for sink, and F for flush)	48
5.1	Comparison of existing techniques to infer usage and power consumption of light fixtures	55
5.2	Three tier data fusion framework used by LightSense to combine data streams. . .	58
5.3	Sample Data Trace from a bedroom light sensor and whole house power meter from 4AM to 4PM in House 2 of our LightSense deployments, along with ground truth fixture events. LightSense eliminates false positive light and power Edges by performing data fusion and matching.	61
5.4	The lpbayes Bayesnet used by LightSense to find the maximum likelihood clusters (top nodes) given the power and light edges (lower nodes).	62
5.5	Optional caption for list of figures	68
5.6	Details of the four homes chosen for deploying LightSense	69
5.7	Energy Costs at \$0.15 per KWh projected over 5 years, for individual light fixtures across all four homes, as computed by LightSense and the ground truth system. LightSense accurately reports the energy costs of the top energy consumers in each home.	70
5.8	Nominal Wattages and Room Locations for the light fixtures in the four houses along with the light intensity increase or decrease observed at our light sensor as the fixtures are switched ON or OFF, respectively. The light fixture numbers here match the light fixture numbers in figure 5.7.	71
5.9	Projected cost savings based on LightSense recommendations and an optimal system with ground truth data.	71
5.10	Edge Detection Recall and Precision for LightSense and its individual inference components. LightSense outperforms the individual components by effectively fusing the noise filtering properties of data fusion and ON-OFF event matching. . .	73
5.11	Relative Total Cost Error for power only assignment and lpbayes power assignment.	75
5.12	Recall Precision tradeoff for LightSense and its individual inference components as the light and power level thresholds are decreased from 500 to 10 in steps 50. . . .	75
6.1	Sensor and sample data used in our resident identification study. Our study used the Go Motion ultrasonic range finder mounted above doorways (a). As users walked beneath the sensor, the range measurements changed (b).	82
6.2	Experimental results for height based identification in the lab. Controlled laboratory experiments indicate low measured error while standing. Mean error while walking is higher due to a natural reduction in height compared to standing erect, and different walking styles. Standard error while walking is very low. The error distribution approximates a log normal distribution.	87
6.3	Identification Accuracy as the number of individuals to be differentiated is increased. Height measurements become less effective for biometric identification as the pool of individuals increases. The heights in our study were easier to differentiate than those of the general population.	88
6.4	Height sensors deployed above doorways in a home.	89

6.5	Sensors used in our height based identification study in homes. Motetrack tags and beacons (left) were used to collect ground truth locations. Motion and magnetic reed switch sensors (middle/right) were used to evaluate STAR.	90
6.6	Accuracy of height based identification and activity model based identification. Height sensors achieve higher accuracy than achieved by existing implementations of the STAR approach that use only activity models and binary sensor use models of residents for identification.	92
6.7	Applicability of height based identification at a nationwide scale. Height sensors are potentially applicable with high accuracy to a large proportion of US elderly households. Weight sensors have potentially higher applicability, but require good sensor design to aesthetically install on the floor.	94
6.8	Potential improvement in height based identification accuracy from using a tracking approach. Our simulation study shows that the history of height measurements collected over the track of a resident through the home potentially improves identification accuracy and applicability of height sensing in US elderly homes.	95
6.9	Impact of tracking on height difference required for accurate identification, and the number of individuals that can be differentiated accurately. A history of readings could potentially decrease the height difference required for accurate identification(a), and increase the number of residents that can be reliably differentiated(b)	96
7.1	Effect of Signal Attenuators on event detection rate (EDR) of Tier II and III events, and on duration accuracy (DA) of Tier 0 events	107
7.2	Effect of random delays on event detection rate (EDR) of Tier II and III events, and on duration accuracy (DA) of Tier 0 events	108
7.3	Effect of signal attenuators, random delays, and fingerprint masking on duration accuracy (DA) and true positive rate (TPR) of Tier II and III events	110
7.4	Effect of period of transmission on node lifetime	111
7.5	Effect of fingerprint masking on event detection rate (EDR) of Tier II and III events	112
7.6	Effect of hybrid schemes on Tier II and Tier III EDR and Tier 0 DA across all homes	113

Chapter 1

Introduction

Wireless sensor systems are increasingly being used to monitor physical phenomena in diverse application domains, such as industrial efficiency and safety monitoring [1], military target tracking and localization [2], outdoor environmental monitoring [3], and urban scale sensing for traffic and environmental applications [4]. In the home environment, a number of exciting medical monitoring and energy conservation applications are enabled by wireless sensor systems that monitor the daily activities of residents [5, 6, 7, 8, 9]. In this thesis, we focus on the problem of **human activity recognition in the home environment**, i.e. inferring the activities of all residents in the home at any time instant.

In the medical monitoring domain, a key application enabled by activity recognition is remote elderly monitoring [10, 7, 8]; these systems monitor elderly residents living independently in their homes, and provide a constant record of residents' daily activity activities to caregivers and relatives of the elderly. Interested stakeholders can then use this daily record to ensure that the residents are able to perform their Activities of Daily Living (ADLs) normally, without fail. Several companies today offer such peace of mind remote monitoring solutions [8, 7], and these systems are already deployed in millions of homes across the world.

In the energy conservation and sustainability domain, information about the daily activities of in-home residents can be used for two distinct purposes - providing better feedback about resource consumption in the home [11, 12, 13, 14], and enabling better control of the appliances in the home to save energy without sacrificing user comfort [15, 5].

Firstly, fine-grained resident activities can be used to infer when individual electrical and water fixtures are being used in the home; this usage data can be combined with smart electric and water meters being installed by utility companies [16, 17, 18], to infer how much energy or water is being used by individual fixtures in the home. In contrast to existing utility bills which only report the total energy or water consumption in the home, smarter utility bills that include information about the resource consumption of individual fixtures in the home have been shown in the literature to be very beneficial in affecting user behavior, and saving energy in the home by as much as 15% [19, 20, 21].

Secondly, activity recognition can also enable smarter appliances that automatically turn on or off, or adjust their operation and energy consumption based on current resident activity. For example, the smart thermostat approach automatically controls the thermostat in an intelligent way to achieve better comfort and energy savings for the in-home residents [15, 5], by leveraging both the instantaneous resident activities, and the long term resident activity patterns. Intelligent home lighting systems that adjust their operation based on resident activities have also been proposed for saving energy in the home [22]. In addition to medical monitoring and energy conservation applications, other applications for resident activity monitoring have been deployed successfully, such as home automation [23], home security [24], and home entertainment applications [25].

1.1 Background

There has been significant work on activity recognition in the home setting, in the context of both medical monitoring and energy conservation applications. Typical activity recognition systems require users to deploy sensors in the home, and apply machine learning algorithms on the sensor data to infer user activities [26, 27, 28]. We argue that existing approaches to activity recognition are **invasive** for the end user, due to three main reasons.

Firstly, most existing activity recognition approaches require significant training from the end user, of the order of several days, to achieve high activity recognition accuracy [26, 28]. This training data is collected through a variety of different techniques, such as using periodic prompts on a smartphone or PDA [29], or using a bluetooth microphone to verbally specify the current activity being performed [26]. Intensive configuration and labeling effort from the end user is a significant

hurdle to the widespread acceptance of home activity recognition systems; end users would like a turn-key activity recognition system that requires little or no configuration.

Secondly, existing activity recognition systems require significant effort from the end user in installing the individual sensors in the home. For example, many activity recognition systems currently in use require anywhere from 40-200 sensors depending on the size of the home [27, 28], and may require on the order of 25-45 man hours in installing the sensors in the home [27, 28, 30]. Some sensors such as magnetic contact reed switches are more difficult to install compared to infrared motion sensors placed easily on a cabinet or a wall. Initial effort in installing and configuring the numerous sensors can be another significant hurdle to widespread acceptance of these systems.

Thirdly, typical activity recognition systems in multi-resident home scenarios inconvenience the end user. In particular, existing approaches to multi-person activity inference requires users to constantly wear inconvenient tags [31, 32, 33], or allow the installation of cameras and microphones [34, 35], which are perceived to be invasive by many end users for long term home monitoring [36]. Multi-person activity recognition is important for both medical monitoring and energy conservation applications. In medical monitoring applications, caregivers would like to know who is currently toileting [7], while in energy conservation applications, rooms could be conditioned to different temperatures or light levels depending on the preferences of a particular user.

1.2 Outline of our activity recognition approach

Our goal is to build convenient, non-invasive activity recognition solutions for end users, by overcoming the three main limitations of existing activity recognition systems. ***Our main hypothesis is that by effectively using data fusion techniques, leveraging the existing smart meter infrastructure in homes, and using only weak biometric sensing, we can build convenient, and accurate home activity recognition solutions for the end user.***

Firstly, we hypothesize that by using effective data fusion techniques, and leveraging cross-home activity classifiers, we can eliminate user effort in training the activity recognition system. Secondly, we observe that many homes are already being fitted with whole-house power and water consumption monitoring units, in the form of smart meters installed by utility companies. We propose a novel data fusion approach to combine these whole house smart meters with a handful of sensors in the home,

to eliminate many individual direct sensors, and reduce user effort in installing direct sensors in the home. Finally, we observe that we can use weak biometric sensors such as height sensors above the doorways of a home to accurately differentiate among the few residents in a home, in order to identify who is performing which activity. Our approach uses only low cost, non-invasive sensors, does not require user training, and uses only a handful of easy to install sensors even in large homes.

We first present our unsupervised activity recognition algorithm [37], which we call **AUTOLABEL**, in chapter 3. We discuss the three tier data fusion approach that combines data from several simple sensors in the home, to automatically infer the room location of the sensors, the objects to which they are attached to, and the daily activities of residents, with little configuration effort from the end user. We present the evaluation of our AUTOLABEL approach in eight diverse home deployments for one week each, to test whether unsupervised activity recognition approaches such as AUTOLABEL can achieve high enough accuracy with little configuration effort from the end user.

After providing a proof of concept for how we can reduce user training effort in chapter 3, we focus on the problem of reducing user effort in installing individual sensors in chapters 4 and 5.

In chapter 4, we present our **WaterSense** approach to eliminate direct sensors on individual water fixtures in the home [38]. Our proposed approach combines a single motion sensor in each bathroom or kitchen, with a whole house smart water meter, to automatically infer when individual water fixtures such as the sink and flush are being used, in each room in the home. To ease end user acceptance, WaterSense uses only low cost, non-invasive motion sensors that are easy to install, and does not require training effort from the end user. Using sensor deployments in two homes, we show that we can effectively disambiguate water fixture usage, by combining both water flow signatures at the water meter, and motion sensor signatures generated as residents in the home use individual water fixtures. Water fixture usage information is important for many activity recognition algorithms, including our unsupervised approach from chapter 3. In addition to monitoring usage, we show that our WaterSense approach accurately infers the total water consumption of individual fixtures in the home.

In chapter 5, we present our **LightSense** approach to eliminate direct sensors on individual light fixtures in the home. Our proposed approach combines a single light sensor in each room, with a whole house smart electricity meter, to automatically infer the number of light fixtures in each room, when they are being used, and how much energy they consume. Similar to WaterSense, LightSense is

unsupervised, and only uses a handful of cheap, easy to install sensors. Based on sensor deployments in four homes, we show that LightSense accurately infers the energy usage of individual light fixtures in the home, which are the highest energy consumer in the low power appliance range. Monitoring light fixture use is also important from a medical perspective, since the usage of certain lights could help differentiate between medically relevant activities, such as sleeping vs. reading, or cooking vs. eating, by inferring the use of bed-side lamps, stove-top lamps, or task lighting near the dining table.

In chapter 6, we present our solution for identifying residents in multi-person homes for activity recognition purposes [39]. We propose using ultrasonic height sensors above doorways for resident identification in homes. Height sensors are convenient, cheap, and easy to install for end users. Height is typically only a weak biometric, but we observe that it is effective for differentiating among the few residents typically seen in home environments. We also show that it is possible to improve the inherent biometric accuracy of height sensors, by aggregating height sensor measurements at multiple doorways in a home, in a data fusion approach. Our approach shows high accuracy in controlled lab experiments with 20 subjects, and high accuracy in multi-resident elderly homes at a nationwide scale, when extrapolating the accuracy results from our lab experiments.

While our study was the first to propose the use of height sensors for resident identification in homes, recent work in the literature has focused on the challenging problem of using ultrasonic height sensors in in-situ environments for the purpose of continuously tracking residents' locations in homes. This recent work shows that the major challenges in using height sensors for tracking are: dealing with false positives and negatives in the height sensor data due to multi-path effects in homes, inferring user direction when walking through a door, and addressing the annoyance caused to pets in the home. Future sensor fusion solutions to the noise problem in height sensors are discussed in chapter 8.

1.3 Privacy considerations

Protecting the privacy of residents in our home monitoring scenario is a significant obligation and design consideration for home activity recognition researchers. ADLs are typically very personal and private, and must be kept secret from third parties. This is particularly true in medical facilities

where ADLs are used to infer medical conditions; these facilities are *obliged* by HIPAA regulations [40] to protect this information.

Our proposed activity recognition approach, similar to other approaches, proposes using *wireless* sensors in the home environment. Wireless technology has clear advantages over wired networking techniques, in particular in enabling convenient and fast installations. For example, over 32 million homes in the US already have wireless security sensors installed on doors and windows, and motion sensors installed inside and/or outside the home [41]. These sensors recognize ADLs of residents for home security and home automation applications, for remote elderly medical monitoring, and also for energy conservation applications.

We make the key observation that our *unsupervised activity recognition* approach introduced in chapter 3 can be used to infer the private ADLs of residents in smart homes, by an adversary *outside the home* snooping only on the wireless fingerprints [42, 43, 44] and transmission timestamps of sensor firings from *inside the home*. Wireless fingerprints serve to uniquely identify the source of timestamped sensor firings, while our unsupervised activity inference approach automatically infers the sensor semantics, and activities performed using our three tier inference algorithm. We call this the *Fingerprint And Timing-based Snooping (FATS)* attack [37].

In chapter 7, we propose a suite of techniques to protect against the FATS attack. For example, we hide packet transmissions from the adversary using signal attenuators, we introduce random delays on transmissions, and we generate spurious transmissions in order to decrease the effectiveness of the FATS attack. We find that each of these techniques has a different tradeoff in terms of privacy protection performance, hardware costs, and application/user costs. Based on the cost-benefit profile of each technique, we present guidelines for applying them most effectively in typical home sensor systems for activity recognition. We evaluate these design guidelines and show how many of these guidelines can be used together in a hybrid fashion to yield very high privacy protection with minimal implementation costs.

1.4 Limitations

Our proposed activity recognition solutions and their evaluation have several limitations that could be addressed in future research.

Firstly, our unsupervised activity recognition approach presented in chapter 3 is evaluated primarily in graduate student and post doctoral employee homes from a single city, and not in diverse geographic locations, or in large family homes and multi-story homes. We have also not evaluated how our algorithm may be extended to address long term deployments of one year or longer, or dense sensor deployments containing more than 100 sensors per home. Since our AUTOLABEL approach is based on the invisible man approach to activity recognition, it can only recognize activities that involve the use of instrumented objects in the home, or simple activities that can be inferred based on time of day, room location, and length of room occupancy inferred using binary occupancy sensors. Our ground truth methodology requires residents and researchers to label the activities post-hoc from the sensor data, and may contain errors in the activity labels due to missing sensor data or poor resident memory; in the future, more accurate ground truth activity labeling systems need to be developed.

Secondly, our WaterSense approach presented in chapter 4 is currently evaluated only in two homes over a short term deployment of one week each. Also, we only consider usage information from sinks and flushes in these homes, since other fixtures are not instrumented for ground truth usage measurement. We combine the data from contact reed switches with the whole house water flow data to estimate the usage and water consumption of individual fixtures in the home; this approach results in errors in the ground truth due to false positives and negatives from the contact switches. As a result, we do not evaluate recall and precision for the fixture usage inferred by WaterSense, but only classification accuracy for true positive water fixture events estimated by the reed switch data and the whole house water flow data.

Thirdly, our LightSense approach presented in chapter 5 could also be improved to address several limitations. LightSense does not report the physical location of the light fixtures in each room. LightSense does not recognize the use of dimmable lights whose intensity changes with every use, nor does it accurately detect light fixtures with low wattage less than 10 W due to high noise at the power meter. In larger rooms, LightSense might require multiple light sensors per room if certain light fixtures do not create any light intensity change at any single sensor location. Another limitation of our current evaluation of LightSense is intermittent data loss from both the light sensors and the whole house power meter during the deployment period, due to power line noise and lossy wireless communication.

Our height sensor based identification approach presented in chapter 6 has several limitations. Firstly, a key limitation of our approach is that our in-situ evaluation of height sensing accuracy does not consider false positives and negatives in individual height sensor events due to the lack of accurate ground truth. False positives and negatives, and height sensor inaccuracy, are possible due to multi-path effects from ultrasonic sensors, especially in (1) rooms with wide doorways that require several adjacent height sensors to achieve sufficient coverage, and (2) adjacent doorways very close to each other. Another key limitation is that we do not develop and evaluate a complete tracking approach to estimate the residents' current room locations based on the height sensor data. An inherent limitation of our approach is that it only identifies residents at doorways, and not inside rooms; such an approach is suitable for room level localization of residents. Finally, since our approach is based on a resident biometric, it cannot be applied in all homes with high accuracy, unlike existing approaches such as tag and track approaches [31, 32, 33].

Finally, in our wireless fingerprint based snooping attack presented in chapter 7, a key limitation is that we do not evaluate the ground truth packet reception rate, and fingerprinting accuracy, for a typical adversarial snooping device deployed outside the home. Another limitation is that we emulate many of our privacy solutions such as radio signal attenuation and fingerprint masking in our evaluation. It would be important in future research to physically implement these solutions to better understand the cost vs privacy tradeoffs involved in these solutions.

1.5 Summary of Contributions

We summarize the five main contributions of this thesis below. In chapters 3 to 7, we list the more detailed contributions of each component of our non-invasive activity recognition approach. Each of our activity recognition components are evaluated through individual, short term home sensor deployments lasting from 7 to 10 days. The evaluation of each of our activity recognition components in real homes constitutes an important contribution in itself.

1. An unsupervised activity recognition algorithm (AUTOLABEL) that requires little user configuration or training effort, and achieves 80-90% activity recognition accuracy, on par with existing supervised approaches that require training over multiple days. Our novel 3 tier AUTOLABEL algorithm is the first to combine unsupervised sensor clustering based on relative

room location, along with cross-home, tiered classification of room locations of sensors, sensor semantics, and daily activities, with little configuration input from the end user.

2. An unsupervised data fusion algorithm (WaterSense) to eliminate direct sensors on water fixtures by combining a single motion sensor in each room with the whole house smart water meter provided by utility companies; WaterSense achieves only a 10% loss in accuracy compared to sensor intensive approaches that require a direct sensor on each water fixture. We design a novel Bayesnet approach as part of WaterSense to fuse ambiguous water flow profiles and motion sensor profiles to accurately identify water fixture usage in homes.
3. An unsupervised data fusion algorithm (LightSense) to eliminate direct sensors on all light fixtures by combining a single light sensor in each room with the whole house smart power meter provided by utility companies; LightSense achieves only about a 10% loss in accuracy compared to sensor intensive approaches that require a direct sensor on each light fixture. As part of LightSense, we design a novel Bayesnet approach to perform data fusion on noisy sensors with low precision to accurately identify physical phenomena that impact multiple sensors.
4. A new sensing approach based on ultrasonic height sensors for convenient resident identification in homes. We show that the inherent biometric accuracy of height sensors can be improved by aggregating multiple height measurements at different doorways in a home. A first pass evaluation shows high identification accuracy in controlled environments, and the potential to achieve 95% identification accuracy in 95% elderly homes in the USA through an extrapolation study.
5. The discovery of a novel wireless snoop attack on smart home activity recognition systems, and low cost privacy solutions to mitigate the adversarial inference accuracy of such an attack from 90% to less than 15%. Our FATS attack is the first to show that only the timestamps and physical characteristics of wireless transmissions can be combined to accurately infer the daily activities of residents in spite of encryption, by learning the relative room locations of sensors over time, and automatically inferring the semantic room locations, sensor types, and resident activities by leveraging the cross-home classifiers from AUTOLABEL.

In chapter 2, we discuss the state of the art activity recognition approaches and their limitations in more detail. In chapters 3 to 7, we present our activity recognition approach, as outlined in section 1.2. In chapter 8, we present the main conclusions drawn from this thesis, and also present the current limitations, and opportunities for future work to address these limitations.

Chapter 2

Background

Several approaches have been proposed in the literature to address the problem of activity recognition in a home environment. A naive approach to activity recognition is to simply employ an observer in the home to manually record the activities of residents in the home. A variant of the manual recording approach is to have the residents themselves self-report their daily activities. Such a manual recording approach is clearly too intrusive for the residents in the home, and is also expensive if a dedicated employee must be paid to interpret and record activities. To ease the burden of self-reporting, experience sampling methods have been explored in the literature, that prompt users periodically to select their current activity from a drop-down list on a personal device such as a smartphone or a PDA [45, 29]. However, these techniques are also too disruptive for end users, and the amount of activity data collected reduces over time as users lose interest in the experience sampling system [46].

To address the limitations of manual activity recording approaches, several researchers have proposed *sensor based approaches* to activity recognition in the home environment. These approaches require users to deploy sensors in the home environment, and then use machine learning techniques on the collected sensor data to automatically report resident activities [28, 27, 26]. Existing activity recognition approaches differ greatly in the types of sensors used, and also the inference techniques used to infer resident activities from the sensor data. The dollar cost of these systems, and the user effort in *installing* the sensor system in the home, depend heavily on the sensor types used in these systems. On the other hand, user effort in *configuring* the sensor system depends heavily on whether the inference techniques used are unsupervised or supervised. We provide an overview of the existing

sensing and inference approaches to activity recognition in the home, and provide a motivation for why our proposed activity recognition system is required to overcome the limitations of these existing approaches.

2.1 Sensing Approaches

2.1.1 Video and Audio based sensing

The earliest activity recognition systems used camera sensors to recognize resident activities in the home. Most of the research was performed by the vision community, which focused on the challenging problem of identifying everyday actions in the home environment, such as stirring a bowl, opening or closing an oven, or using a toaster [47, 48]. While camera based systems are promising, a key hurdle to their widespread acceptance is that end users perceive cameras to be too invasive for long term monitoring in a home environment. In spite of guarantees about how the data is stored or interpreted, user studies conducted in the literature found cameras, along with microphone sensors, to be most objectionable to end users for long term activity monitoring [36, 49].

2.1.2 Invisible man approach to activity sensing

To ease end user concerns about continuous audio or video recording, researchers at MIT and Intel proposed the pioneering invisible man approach to activity recognition in the home. The key breakthrough here was to replace invasive cameras or audio sensors, with simple, binary sensors on objects used in everyday activities, such as the stove, microwave, toilet, shower, along with simple binary motion detectors in rooms to detect user occupancy [28, 27]. The term “invisible man” was coined to denote the fact that the residents themselves were invisible to the sensor system, and were not being monitored using video cameras; rather, only their interactions with everyday objects in the home were monitored for activity recognition purposes. For example, the use of appliances such as the stove or the microwave could be used to infer cooking, while the use of water fixtures such as the flush and sink could be used to infer toileting activities. By being able to directly recognize interaction with everyday objects, the invisible man approach also avoided the challenging problem of recognizing objects using video cameras in a dynamic environment such as the home. Several user studies, and reports from commercial home monitoring companies, confirmed that this invisible

man approach was most acceptable to end users for activity recognition in the home environment [27, 50, 49, 36]. One key limitation of the invisible man approach is that it is only able to recognize activities that involve interaction with objects in the home; thus diverse physical activities that do not involve interaction with everyday objects such as jumping, sitting, exercising, or lying on the floor, cannot be recognized using simple binary sensors alone. To recognize such activities, more invasive camera sensors [47] or inconvenient, wearable accelerometer sensors may be used [51]. However, the invisible man approach is able to recognize the majority of activities and room locations relevant to many medical monitoring and energy conservation applications, using only non-invasive, binary sensors.

The invisible man approach was implemented using two main sensing modalities, which presented different tradeoffs in terms of usability, and sensor cost and installation effort.

Wearable sensing

Researchers at Intel proposed the iBracelet and iGlove wearable system for monitoring everyday interactions with objects in the home [52, 53]. The wearable bracelet was equipped with an RFID reader with a small read range, while everyday objects were tagged with cheap RFID tags, costing about 10 cents each. When a user interacts with an object such as a kettle, or the stove, the wearable RFID reader reads the unique RFID tag on the object, to recognize which object the user has interacted with [53, 54]. Thus, at the cost of a more expensive wearable device, Intel researchers were able to drastically reduce the cost of individual sensors attached to objects to monitor their usage. However, the main drawback of wearable sensing approaches is that, for long term home monitoring applications such as peace of mind monitoring [7] and energy conservation [5], users are uncomfortable wearing or carrying a sensor device all the time. Even if residents do agree to a wearable sensor, we observed in our own deployments [30] that users often forget to carry the tags, resulting in lost activity data.

Non-invasive direct sensing

To overcome the inconvenience caused by continuous wearable sensors for long term home monitoring, Tapia et al at MIT developed an activity recognition system using simple and ubiquitous self-contained sensors in the home environment [28, 27]. In particular, magnetic contact reed switch

sensors were used to monitor the usage of everyday objects such as the sink, the flush, the shower, the microwave, and the fridge. By combining these sensors with non-invasive, binary infrared motion detectors in the home, Tapia et al achieved good results in recognizing a number of resident activities relevant to both medical monitoring and energy conservation applications [28]. However, a significant drawback to this sensing approach was that users still had to spend a large number of man-hours installing the sensors in the home. For example, in many of our home deployments, an average of 20 man hours per home was spent just on installing direct sensors to monitor the use of light switches in the home. The sensors also required constant monitoring for failures, and were not aesthetically pleasing, since they were installed conspicuously on almost every object the user interacted with in the home [30].

Emerging single point sensing

To overcome the limitations of non-invasive direct sensing, researchers at Georgia Tech and the University of Washington, propose the use of single point sensors on the home's electrical, water, and gas lines [55, 56, 57, 58, 59], to effectively eliminate direct sensors on individual water fixtures such as the sink and the flush, and individual electrical fixtures such as light fixtures, the stove and the microwave. These approaches use differences in the noise transients created on the home's electrical or water lines as fixtures are used, to identify the usage of individual water and electrical fixtures for use by the invisible man approach. However, these approaches require significant configuration and training effort from the end user, high cost sensors, and have not yet been evaluated in in-situ home environments. We discuss single point sensing approaches in more detail in chapters 4 and 5, when we discuss our non-invasive approaches to eliminate direct sensors on electrical and water fixtures in the home by leveraging the smart metering infrastructure.

2.2 Inference approaches

Irrespective of the sensing approach used to infer when individual objects are used in the invisible man approach, a significant proportion of research on activity recognition has focused on the learning techniques used to infer resident activities from the sensor data. We classify the learning approaches

into two broad categories depending on the amount of user input required to configure or train the activity recognition algorithm.

2.2.1 Supervised learning

Supervised learning approaches require users to label their current activities over several days or weeks [28, 26]. Based on the labeled data, optimal classifiers are created to translate the input sensor data to output activities, by leveraging several existing supervised classification techniques, such as Bayesian classifiers, hidden Markov models, support vector machines, and decision trees. Although supervised classification approaches are an improvement over self-reporting systems, they still require residents to provide a constant record of their daily activities over several days or weeks, using devices such as a Bluetooth microphone [26], or an experience sampling technique using a SmartPhone or a PDA [28, 29]. Such high configuration effort appears to be a significant hurdle to the widespread acceptance of activity recognition systems in the home. Ideally, users would like a turn-key activity recognition system.

2.2.2 Unsupervised learning

Unsupervised learning approaches do not require continuous user effort in training the activity recognition system. The most significant unsupervised activity recognition approach was developed by Wyatt et al in 2005 [54]. This approach requires an RFID reader bracelet to obtain object use information from tagged objects; this approach then uses automatically mined activity models from the web to recognize activities, by viewing activities as a sequence of natural language terms, where the individual terms represent the use of a particular object such as the sink, tea kettle, or flush. A significant drawback of this approach is that it requires users to label each sensor in the home with a textual semantic label that ideally matches text corpora in the web describing the object. Thus, it transfers user configuration effort from labeling daily activities, to labeling the numerous dense sensors in the home. With this approach, the researchers achieved only about 42% recognition accuracy, well above the 3.8% baseline, but still too low for many activity recognition applications [7, 5].

To overcome the limitations of these inference approaches, we propose the AUTOLABEL unsupervised activity recognition approach in chapter 3 [37]. Our approach neither requires users to label individual sensors for everyday objects, nor requires users to label daily activities. Our AUTOLABEL approach automatically infers the location and semantics of the objects to which individual sensors are attached to, and also the daily activities of residents in the home.

In chapters 4 and 5, we show how to improve over existing sensing approaches to provide the underlying input to our AUTOLABEL approach, by *eliminating many of the direct sensors used in our proof of concept activity recognition evaluation in chapter 3* [38]. Finally, an important concern for home activity recognition systems is the problem of differentiating activities of multiple residents in the home, by identifying their location without using invasive wearable sensing or camera sensors. In chapter 6, we propose a non-invasive convenient solution for resident identification in homes, that shows high accuracy in initial experiments in the lab [39]. In chapter 7, we present a novel wireless snooping based privacy attack on resident activity recognition systems, and propose solutions to protect against such an attack in future smart homes [37].

Chapter 3

Reducing user configuration effort

To reduce user effort in training and configuring the activity recognition system, we develop an unsupervised activity recognition algorithm, that neither requires users to provide semantic labels for individual sensors in the home [54], nor requires effort in labeling daily activities for the purpose of supervised classification [28, 26]. We call our activity recognition algorithm as AUTOLABEL [37]. Our AUTOLABEL algorithm uses a tiered inference approach that effectively performs data fusion between the multiple sensor streams to automatically infer the room location and semantics of each individual sensor, and leverages cross-home activity classifiers to accurately recognize daily activities of residents without any training effort from the resident.

We empirically evaluate our AUTOLABEL approach by deploying the same non-invasive direct sensors used in the invisible man approach [28, 26], in eight different homes. We deployed low cost, non-invasive sensors on salient objects, including sinks, stoves, showers, and doors, and collected data for a week or more in each home. We tested a diverse range of homes, with different floor layouts, different sets of deployed sensors, and occupied by both single and multiple residents from different age groups and lifestyles. We used individual direct sensors mainly to provide a proof of concept evaluation of our AUTOLABEL algorithm. In a realistic activity recognition system, many of these direct sensors will be eliminated by using our low cost sensing approaches presented in chapters 4 and 5, that leverage the smart metering infrastructure in homes. We observe that our AUTOLABEL algorithm is surprisingly robust to the diversity of homes, people and sensed objects in our deployments, and can infer detailed resident activities with high accuracy.

The rest of this chapter is organized as follows. In section 3.1, we present the design of the AUTOLABEL algorithm, and illustrate its operation by using example data from one of our home deployments. In section 3.2, we present the sensor deployments in eight diverse homes and the experimental setup, to evaluate our AUTOLABEL approach. Finally, in section 3.3, we present the activity recognition accuracy of our AUTOLABEL approach in the eight homes.

3.1 AUTOLABEL design

We now provide an overview of the activity inference algorithm. Figure 3.1 shows the inference tiers used in the AUTOLABEL approach, to transform the input sensor streams into activity streams, without any knowledge about the location or semantics of each individual sensor stream, or any activity labels for training. To better illustrate the AUTOLABEL algorithm, we depict its concrete operations on real data in Figure 3.2; the various black lines in figure 3.2 depict individual sensor streams in the home, while the vertical spikes indicate the timestamps when these individual sensors fired.

Firstly, figure 3.1 shows a **Tier 0**, where we can only detect very general activities such as home occupancy or sleeping, from just the timestamps of sensor firings. In **Tier I**, we use the unique sensor identifiers for each sensor firing, depicted by the black lines in Figure 3.2. Then, these sensor identifiers are grouped into *sensor clusters* corresponding to rooms in the home based on similarities in their firing times, shown by the cluster labels on the left of Figure 3.2. The approximate number of residents in the home can be inferred by observing simultaneous activity in multiple rooms, as shown in Figure 3.1. In **Tier II**, specific features are first extracted from the combined firings of all devices in a spatial cluster, denoted by the red lines in Figure 3.2 beneath each cluster. These features are passed to a classifier that identifies each room as either a kitchen, bathroom etc. Figure 3.1 shows that these room labels can be used to infer the number of times the residents visit, for example, the kitchen or bathroom each day. Finally, in **Tier III**, another classifier is used to determine the likelihood of a sensor being the motion sensor, stove sensor etc. This information can be used to recognize subtly different activities, namely cooking hot and cold food or showering, toileting, and grooming, as shown in figure 3.1. In the rest of this section, we describe in more detail, the design of the four tiers in our AUTOLABEL algorithm.

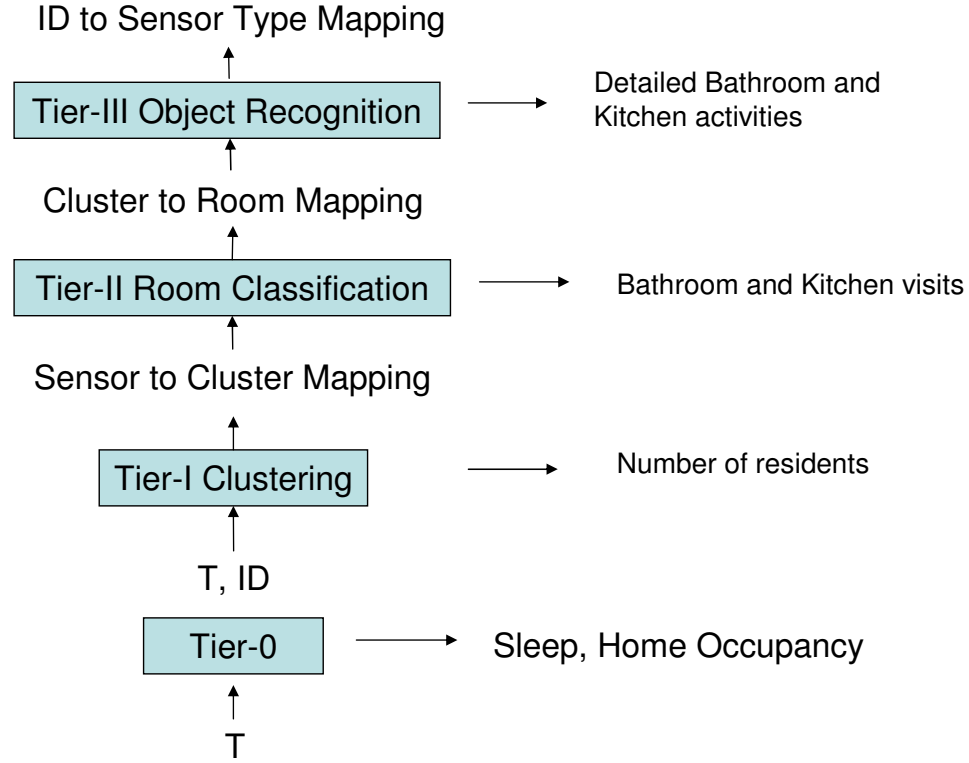


Figure 3.1: The AUTOLABEL Activity Recognition Algorithm that uses just the timestamp T , and unique source identifier ID , of each binary sensor firing to automatically infer daily activities of residents without any configuration or training effort.

3.1.1 Tier-0: General Activity Detection

Firstly, we observe that we can identify long term, coarse grained activities such as home occupancy, and sleeping, by just using simple features such as time of day, and the presence or absence of sensor firings from *all* motion detectors in the home. In Tier 0, we implement a simple algorithm that identifies silent periods during the day as *away* events, silent periods during the night as *sleep* events, and all active periods to be *home* events. To reduce the effect of spurious sensor activity, we use a threshold of at least 4 sensor firings per hour to indicate home presence. In multi-person homes, only the aggregate activity can be inferred, such as when *everyone* in the home was sleeping or not. In Tier 0, we use only the timestamps of the sensor firings. In the subsequent inference tiers, we show that knowledge of sensor identifiers *in combination* with transmission timestamps allows us to infer much more detailed activities of residents in the home.




Figure 3.2: Sample data from one of our home sensor deployments. The textual labels and logical groupings are derived from the raw sensor data (black lines) by our AUTOLABEL algorithm.

3.1.2 Tier-I Clustering

The goal of Tier I is to identify which sensors are in the same rooms. It does this by assuming that sensors in the same room fire at similar times due to human activity in the room. Thus, we use a *temporal distance* calculated between the transmission patterns of each pair of sensors to cluster

together those sensors that have small distances to each other. In chapter 4, we will show how to adapt this room clustering algorithm to handle more complex sensor types, such as using a single water meter to sense individual water fixtures in the home.

We denote the set of all devices (identified using their unique identifiers) to be ID , and the vector of all transmission timestamps from each device $i \in ID$ to be T_i . We use bracket notation to index into vectors, so the k th timestamp from node i is referred to with $T_i[k]$. The Tier I clustering algorithm is then defined as:

```

for all  $i, j \in ID$ 
  for  $h = 1$  to  $length(T_i)$ 
     $dist_{ij}[h] = \infty$ 
    for  $k = 1$  to  $length(T_j)$ 
       $dist_{hk} = |T_i[h] - T_j[k]|$ 
      if  $dist_{hk} < dist_{ij}[h]$ 
         $dist_{ij}[h] = dist_{hk}$ 
     $D_{ij} = \min(\text{median}(dist_{ij}), \text{median}(dist_{ji}))$ 
 $D' = \text{SPDIST}(D)$ 
 $F = \text{CMDS}(D')$ 
CLUSTER = k-means( $F, k$ )

```

For each pair of devices i and j , we first compute the difference in time between each transmission of i and the *closest* transmission of j , creating a difference vector $dist_{ij}$ with length $|T_i|$. We then calculate the **temporal distance** D_{ij} between devices i and j to be the minimum of the median of the time difference vectors $dist_{ij}$ and $dist_{ji}$. These distances are stored as elements of the symmetric distance matrix D of size $|ID| \times |ID|$, which holds the distances between each pair of transmitters.

Not all devices in the same room will be temporally correlated; for example, the dishwasher may have a large temporal distance from the refrigerator because they are never used together, even if both devices have small temporal distances to something else like the sink or a motion sensor. Therefore, we use Dijkstra's shortest path algorithm to convert the distance matrix D to a new matrix D' of metric distances by replacing each distance D_{ij} by the shortest path distance SP_{ij} through D . We then use classical non-parametric multi-dimensional scaling (CMDS), to convert the distance matrix D' into positions of the $|ID|$ sensors in $|ID|$ -dimensional space. Finally, we use the k -means

clustering algorithm to cluster together sensors that are temporally correlated, arriving at a mapping $CLUSTER : ID \rightarrow C$, from the source identity $i \in ID$ of every device to one of the $k = |C|$ clusters; our assumption is that each cluster consists of sensors from the same room, and there exists at least one cluster for each room in the home. Takada et al [60] describe a related algorithm to group co-located sensors together based on time series data; however, the underlying sensors used, the distance matrix computation and the clustering algorithm used there are different. A key requirement of the Tier I clustering algorithm is that the end user has to manually input the number of rooms or k for the k-means clustering algorithm. In the Tier-I evaluation in section ??, we show that the algorithm is robust to incorrect values of k within two rooms of the actual number of rooms in the home, and is able to correctly cluster the majority of sensors in the home.

3.1.3 Tier-II Room Classification

The goal of Tier II is to identify the function of each room as a bathroom, kitchen, bedroom, or living room, without requiring any configuration input from the end user. This tier makes two assumptions: (i) different houses have similar rooms (ii) similar rooms across homes can be identified using specific features of room usage. To identify the function of each room, this tier passes features computed for each room based on the the overall sensor transmissions from the room to a bi-partite matching based classifier. From assumptions (i) and (ii), the bi-partite matching based classifier is created using training data from a small set of homes, and then is used repeatedly in multiple test homes to label rooms; thus, the user is relieved of the need to explicitly enter room labels into the activity recognition system. We call this type of classification approach as a **cross-home classification** approach.

Once the devices are clustered, we generate an overall series of timestamps T_c of all transmission timestamps from all sensors in sensor cluster $c \in C$. We also generate **temporal activity clusters**, used in the features below, by using the *db-scan* algorithm to cluster the timestamps in T_c . Each temporal activity cluster forms a continuous temporal block from T_c with a relatively high density of sensor firings. *db-scan* [61] performs well here because it automatically leaves out outliers and computes high-density clusters unlike *k-means*. We then generate a number of features for every room cluster c from the series T_c listed below.

- the number of transmissions per day from the room
- the median inter-transmissions time within a room

- the median length of temporal activity clusters
- the total number of transmissions during the day
- the total number of transmissions during the night
- the cluster to transmit first after long silence periods
- a histogram of transmission, with four hour granularity

We use these features to create a feature vector F_r for every room r . To classify rooms in a test home, we set apart a small number of other homes from our deployment to provide *training data*. We define R to be all possible room labels (e.g., bedroom, kitchen, bathroom, etc) and create a single feature vector F_r for every label $r \in R$ by averaging the feature vectors of all rooms in the training data with the same label. Then, to label the rooms in a test home, we compute a *min-cost bipartite* matching between the feature vectors $F_c : c \in C$ from the clusters of that home and the training feature vectors $F_r : r \in R$. We define the cost of a match between two feature vectors to be the sum of the *Euclidean distances* between corresponding individual features.

The resulting bipartite matching represents the room labeling. Unlike a conventional classifier, the matching process allows us to enforce mutual exclusivity of room labeling; for e.g., a home with three rooms cannot have two bathrooms or three kitchens. When necessary, we can allow for multiple rooms with the same label in a home by simply including extra copies that room in the set of room labels. In our experiments, we observed that all rooms were correctly labeled across both single and multi-person homes. Using room labels from this tier, we infer the timing and duration of individual room visits, such as *bathroom and kitchen visits* by simply using the temporal activity clusters that occur in each room classified as the bathroom or kitchen.

3.1.4 Tier-III Sensor and Activity Classification

The goal of Tier III is to identify activities in the home, such as cooking, showering, or toileting. This is a two step process. In step 1, we classify sensors to obtain a probability distribution for each unknown sensor indicating the likelihood of its matching with known sensors such as the stove, shower etc. The matching probability M_s for sensor s is indexed by the known sensor types; for example, $M_s[\text{stove}]$ indicates the probability that sensor s is the *stove* sensor. In step 2, we use these probability distribution vectors to classify activities based on which sensors are likely to be active during an activity. Similar to Tier II, we use a cross-home classification approach to identify activities, since we observe that across homes, similar sensors are involved in the same activities;

for example, the flush sensor for toileting, the stove sensor for cooking, etc. Thus, end users are relieved of the burden of having to constantly provide training data to the activity recognition system over multiple days. Interestingly, we do not require the sensor classification to be 100% accurate and show here how inaccurate probability distributions for sensor labels can be used to accurately classify activities.

Sensor Classification

To obtain the probability distribution over known sensor labels for each unknown sensor stream, we first calculate a feature vector for each sensor, using features similar to those described in Tier II, obtained from temporal characteristics of the sensor firings. We then classify each sensor by passing each feature vector through a standard *linear discriminant analysis* (LDA) classifier, trained again using a cross-home classification approach. In other words, for each of the eight homes, we constructed this classifier using training data obtained from a subset of the remaining seven homes with hand-labeled sensor types. We do not assume that all houses have exactly the same set of sensors, but we do assume that all types of sensors in the test home have been observed at least once in a training home.

We train a separate classifier for each room in a home: a sensor from the kitchen cluster in Tier II will not be classified as a bathroom sensor. Thus, the room classification results from Tier II are used to improve the sensor classification results in Tier III. The output of this classification procedure is the mapping vector M_s for each sensor s in a home. In our experiments, the mapping vectors were accurate for bathroom sensors, but were not very accurate for kitchen sensors, where activities like cooking produced similar features in several sensors, such as the stove and cabinet sensors. These objects are often either misclassified as each other or partially classified as multiple objects. However, these incorrect classifications can still be used to recognize activities in Tier III.

Activity Recognition

To recognize activities, we calculate a feature vector for every *temporal activity cluster* in every device cluster. These feature vectors include (i) start time, (ii) duration, and (iii) the total number of times that each known sensor type transmits. Feature (iii) is obtained by adding the mapping vectors for every sensor firing in the activity cluster. If a device is partially classified as multiple

types of sensors, partial counts are maintained in the feature vector. This ensures that sensors used in the same activity (such as the pantry and stove used in cooking) that are misclassified or partially classified as each other do not affect the overall counts of known sensor firings in feature (iii). We then classify each *temporal activity cluster* as an activity using an LDA classifier that was trained on other homes with hand-labeled activity labels. We used this approach to recognize *showering*, *washing/grooming*, and *toileting* activities in the bathroom, and to recognize *cooking hot food* and *preparing cold food* in the kitchen.

Thus, our three tier AUTOLABEL inference algorithm automatically infers the sensor semantics and daily activities of residents in the home, without requiring any configuration input from the end user. Our approach is a significant improvement over existing approaches that require significant configuration or training effort from the end user in the home.

3.2 Experimental setup

3.2.1 Home Sensor Deployments

We empirically evaluate our AUTOLABEL approach by collecting real sensor data from homes containing off-the-shelf wireless X10 motion sensors and contact sensors. We used an X10 receiver inside the home to record the timestamp and device ID for each sensor firing message. Motion sensors were placed in each room and contact switches were placed on doors, sinks, toilets, showers, refrigerators, stoves, and cabinets, some of which are shown in Figure 3.3. These sensors are the same types of non-invasive low cost sensors used in the invisible man approach. Even though we use direct sensors for our proof of concept evaluation here, we show how to eliminate many of these direct sensors on water fixtures and electrical fixtures in chapters 4 and 5, using smart meters in homes.

Data was collected in each home for seven days or more. To be sure that our evaluation is not specific to a particular home or type of home, we deployed the X10 devices in eight diverse homes. All homes had different floor plans, a different number of sensors, different items being monitored, and a different number of residents with diverse age groups and occupations. Some homes had three male graduate student residents while other homes had a married couple and mother-in-law. Layouts varied from studios to two-bedroom apartments. A summary of the diversity of homes is shown in

table 3.1. Hereafter, we label the single person homes A through D and the multi-person homes E through H.




Figure 3.3: Examples of the sensors we deployed to evaluate AUTOLABEL.

Home	Number of Sensors	Number of Rooms	Number of people	Number of firings
A	13	3	1	5888
B	22	5	1	3074
C	16	4	1	4020
D	16	4	1	10326
E	12	4	3	14340
F	15	5	3	20534
G	14	4	2	10964
H	16	5	2	16571

Table 3.1: Characteristics of the eight homes used in our AUTOLABEL deployments

3.2.2 Evaluation Metrics for Activity Recognition

Each tier of our inference algorithm will produce a set of resident activity time intervals I that are defined by a start time and an end time. We also produce a set of *ground truth* activity intervals \hat{I} by hand-labeling the sensor data after data collection is over. We use a min-cost bipartite matching algorithm to pair each interval in I with an interval in \hat{I} . Then, we use three metrics to quantify the correctness of the inference algorithm:

1. **Event Detection Rate (EDR):** the percentage of intervals in \hat{I} that were mapped to some interval in I .
2. **True Positive Rate (TPR):** the percentage of intervals in I that were mapped to a real event in \hat{I}
3. **Duration Accuracy (DA):** the absolute difference in duration between events in I and their matched events in \hat{I}

A EDR value of 60% would be produced if 10 cooking events occurred but only 6 were detected. A TPR value of 60% would be produced if the adversary detected 10 cooking events of which 4 were false alarms. By measuring both EDR and TPR, we ensure that the FATS algorithm does not achieve a high EDR simply by generating a large number of spurious events. A DA of 60% would be produced if a cooking event takes 100 minutes, and the algorithm indicates an event of either 60 minutes or 140 minutes.

3.3 Evaluation

In this section, we present the activity recognition accuracy of each tier in our AUTOLABEL approach using the evaluation metrics described above. Firstly, we note that sleep and home event detection has 85 to 100% duration accuracy (DA) in both single and multi-person homes. Also, as shown in Figure 3.5, we observe that for these two activities, the event detection and true positive rates, are also very high, indicating that we detect coarse grained home occupancy and sleep occupancy very accurately by just using the timestamps of sensor firings in the home.


3.3.1 Tier-I Evaluation

In Tier I, unlike Tiers II and III, no resident activities are detected. Rather, only sensor clusters corresponding to rooms in the home are created. However, the activity recognition accuracy of activities identified in Tiers II and III, depend on how accurately sensors in the same room are grouped together in Tier I. Therefore, we first evaluate the accuracy of our sensor clustering approach.


To compute clustering accuracy, we first compute a maximal min cost bipartite mapping between the set C of computed clusters and the set \hat{C} of true room clusters obtained from the deployment plan. We define a device $i \in ID$ to be *correctly clustered* if i 's computed cluster $c \in C$ is mapped to a room cluster $r \in \hat{C}$ such that device i is actually in room r . We then define **clustering accuracy** to be the proportion of devices that are correctly clustered.

Figure 3.4 shows the clustering accuracy across homes for multiple values of k , and the *non-parametric db-scan clustering algorithm* [61]. Maximal clustering accuracy is achieved when the value of k matches the number of rooms in each home. However, *our overall inference algorithm does not require the exact number k of rooms in the home and is robust to incorrect, larger values k .*

The parameterless db-scan algorithm performs poorly in several homes where all devices are highly correlated temporally. Thus, we need a parameterized clustering technique like k-means with large enough k to enforce a reasonable partitioning of the sensors into clusters corresponding to rooms in the home. It is important to note in figure 3.4b that clustering accuracy remains high in *multi-person homes* in spite of simultaneous activity in different rooms; this is because we still get sensors firing in the same room most of the time.



(a) Single Person Homes




(b) Multi-Person Homes

Figure 3.4: Accuracy of clustering sensors into rooms in AUTOLABEL, using both *dbscan* and *k*-means with 4 different values of k .


3.3.2 Tier-II Evaluation

In Tier II, from the room labels attached to sensor clusters, and the temporal activity clusters inferred for each room, we infer the room visits of residents in the home to rooms such as the bathroom and the kitchen. Figure 3.5 shows that AUTOLABEL infers these room visits, with 95-100% accuracy across both single and multi-person homes using our room classification algorithm. Of course, in

multi-person homes, the variables represent how many times *all* residents in the home visited rooms such as the bathroom or kitchen in total.



(a) Single Person Homes



(b) Multi-Person Homes

Figure 3.5: **Overall Inference Accuracy** of AUTOLABEL on events detected in Tier 0, Tier II, and Tier III, using all three metrics: EDR, TPR, and DA.

3.3.3 Tier-III Evaluation

In Tier III, our AUTOLABEL approach infers a probability distribution over sensor labels for each sensor, and also the daily activities of residents in the home. Though we do not show the sensor label distributions here, figure 3.5 shows the activity recognition results for Tier III. In single person homes, the average accuracy of AUTOLABEL in inferring detailed activities is around 80%, while the accuracy in multi-person homes is lower, but still well above the baseline of random guessing. The lower accuracy in multi-person homes is due to spatial clustering errors introduced by simultaneous activity in many rooms; for example, in home E, an incorrect clustering of the shower sensor with the living room prevented us from detecting showering events. To address the clustering and sensor

classification errors in multi-person homes, we note that in realistic deployments as discussed in chapters 4 and 5, prior knowledge about each water or electrical fixture can be used to identify its semantics based on whole house smart meters, and this information can feed back into room clustering decisions. Also, our approach for identifying multiple residents in chapter 6, can also be used to eliminate some of the spatial clustering errors introduced in our sensor deployments discussed above.

3.4 Limitations

Even though our AUTOLABEL approach achieves high activity recognition accuracy in our sample set of homes, there are several limitations with both the approach and the evaluation that could be improved as part of future research.

A key limitation of our AUTOLABEL evaluation is that it is largely carried out in relatively similar graduate student or post-doctoral employee homes within the same city. We did not encounter large families with young children, or large multi-story homes with diverse room functions. Regional or cultural variations may result in diverse usage patterns, and our current AUTOLABEL algorithm may need to be improved to handle such variations. For example, the end user could manually input their lifestyle or family size from a pre-defined list, each with its own training set for our AUTOLABEL algorithm; such lifestyle information could also potentially be learned automatically by examining which training set maximizes the likelihood of observing the home sensor data. In our sample set of deployments, we do not consider the challenge of how to provide information back to the user regarding the level of uncertainty in the room locations, sensor types, and activities recognized by AUTOLABEL; such feedback could potentially allow the user to more efficiently correct errors in the unsupervised sensor labeling producing by AUTOLABEL.

Another limitation is that the AUTOLABEL approach is evaluated only using short term deployments of 7 days each. We have not addressed the problem of adapting our activity recognition algorithm automatically to account for changing object usage patterns over the long term when performing daily activities.

Since our AUTOLABEL approach uses the same sensing approach as the invisible man approach to activity recognition, it shares its key limitation: AUTOLABEL only recognizes activities that

involve the use of instrumented objects in the home, or simple activities that can be inferred based on time of day, room location, and length of room occupancy inferred using binary occupancy sensors. Diverse physical activities such as jumping, or exercising that do not involve interaction with objects in the home, cannot be recognized with high accuracy by our approach. However, to infer such detailed activities, state of the art techniques use invasive sensors such as camera sensors [47], or inconvenient inertial sensors attached to the body [51]. In contrast, our AUTOLABEL approach uses only non-invasive sensors, but cannot infer detailed activities that do not involve interaction with objects in the home.

Additionally, our home sensor deployments contained about 15-20 sensors in each home depending on the size of the home; in contrast, some sensor deployments of the invisible man approach contained a high density of sensors with around 80-100 sensors per home, including sensors on objects such as the brush, tea pot, and cooking pan [28, 27, 54]. We have not evaluated the accuracy of our sensor labeling and activity recognition approach, when using dense sensors to infer fine-grained activities such as making tea, or brushing teeth. In the future, it would be interesting to evaluate if our high sensor labeling accuracy extends as the density of sensing increasing, and which sensors, if any, are incorrectly labeled by the AUTOLABEL approach.

Although our AUTOLABEL approach is unsupervised, and can automatically infer sensor labels and daily activities, it does require the user to input the number of rooms for our Tier I clustering algorithm; an interesting direction for future research is to explore clustering approaches that eliminate the need to enter the number of rooms. As mentioned above, in the future, the user may also interact with the system by entering some minimal lifestyle information, or interact with the system based on the uncertainty about sensor labeling and activity data to correct erroneous inference. Our unsupervised AUTOLABEL approach achieves high recognition accuracy, of 85-90% that matches the recognition accuracy of supervised approaches evaluated on similar home deployments in the literature [26]. In contrast, AUTOLABEL does not require users to label their daily activities over multiple days, or manually enter the location and semantics of each sensor unlike other unsupervised approaches [54]. However, a limitation of our evaluation is that state of the art supervised learning algorithms have not been implemented on our home sensor data to compute the potential improvement in accuracy that supervised approaches could achieve on our sample home sensor data.

Finally, another limitation of our evaluation is that we manually label the activity data post-hoc in an indirect manner using researchers and home residents to interpret the sensor data; such an approach may not have the same accuracy as a direct, but inconvenient approach where the end user personally labels the activities while performing them [26]. Convenient ground truth annotation for activity recognition research is a challenging and expensive problem. In the future, it would be important to develop more rigorous techniques to label the activity data automatically with high accuracy for ground truth purposes, or adopt more cost and effort intensive approaches as used by the placeLab deployment, where subjects are paid to label the activity data post-hoc based on a video camera placed in the home that records resident activities [27].

3.5 Conclusions

We note that the high proportion of activities that are correctly classified in Tier III indicate that the simple assumptions made by our AUTOLABEL algorithm in Tiers I, II, and III appear to be true across our sample set of homes, and that this type of tiered activity recognition approach can be used to accurately infer activity information about the residents of a home. Our three tier inference approach is designed to first infer the room location and semantics of each *binary sensor* used as part of the invisible man approach, and then use this information in a cross-home activity classifier. Thus, unlike existing approaches, the user neither has to enter configuration input about the sensors in the home [54], nor label daily activities unlike supervised approaches [26, 28]. Our three tier approach is specifically designed to achieve high accuracy in inferring the room locations and fixture types of each sensor. We observed that first labeling the rooms in Tier II based on their temporal usage patterns was highly accurate. Also, a key advantage of *first* determining the room location of sensors is that it narrows down the potential space of fixtures that can be present in the room; this improves the classification accuracy of determining the correct fixture types greatly, since a flush and a stove cannot be confused since they are unlikely to be in the same room.

Our three tier approach may be modified to accommodate emerging sensor types such as *single point sensors* to replace some of the binary sensors on water fixtures and electrical fixtures; we discuss these extensions in the next two chapters. For example, in our current approach, detecting usage events from individual objects in the home such as a sink or a flush is straightforward, since

we use a binary event sensor per fixture. However, when using single point sensors, a separate *edge detection tier* is needed to detect unique events from the single sensor on the home's power line or water input.

In the AUTOLABEL algorithm presented, grouping sensors into rooms in Tier I is performed by clustering distinct sensors based on a unique temporal distance metric between every pair of sensors. However, when using single point sensors, there are no unique sensors to represent the usage of unique fixtures in the home; rather, more sophisticated Bayesian algorithms and filters are required to cluster individual noisy events on the power or water line into rooms, by combining data from the single point sensors and a few distributed sensors in the home. Finally, the signatures generated on the water or power line for single point sensors may contain information useful in identifying the type of sensor currently inferred in Tier III of the AUTOLABEL algorithm; a Bayesian approach may be designed to leverage this information, and jointly infer the relative room location, room label, and sensor label, instead of the tiered approach presented above.

Chapter 4

Reducing sensor installation effort to monitor water fixture usage

The AUTOLABEL approach presented in the previous chapter eliminates user effort in training and configuring the activity recognition system. However, practical activity recognition systems must also aim to reduce user effort in installing individual sensors on objects used in the home. In particular, the advantage of not having to configure each sensor by using our AUTOLABEL algorithm may be lost, if the user has to physically install sophisticated sensors on each fixture. The contact reed switches used in the previous chapter to monitor water fixture usage are low cost and non-invasive; however, these magnetic reed switches are quite difficult to install for a novice, and are also inconvenient for everyday usage, since the aesthetics of water fixtures are affected by the presence of a conspicuous sensor on each faucet, flush or shower [55, 30]. In addition to being difficult to install, magnetic reed switches only report when each water fixture is being used, and not how much water is consumed at each fixture, which is an important need for sustainability applications. In this chapter, we focus on the problem of reducing sensor installation effort to monitor the usage and water consumption of water fixtures in the home.

As seen in the previous chapter, water fixture monitoring is essential for recognizing several important medically relevant activities of daily living, such as cooking, washing, toileting, and showering [28, 26, 37]. In addition to medical activity recognition applications, water fixture monitoring has important implications for sustainability and resource conservation applications. The

world's usable water supply is decreasing at a faster rate than it can be replenished. Household water conservation is important to ensure sustainability of fresh water reserves, to save energy from water treatment and distribution, and to prevent fresh water habitats from being affected through excessive water use [62]. Residents have a number of practical options to conserve water, ranging from replacing high flow toilets and showers with low flow replacements, to reducing water usage for daily activities such as brushing teeth or washing dishes. To make informed decisions that maximize water savings, households first need a detailed understanding of *how much water* is used by each appliance and water fixture in the home.

A simple alternative to magnetic contact reed switches is to install a water flow meter at every fixture to monitor both usage and water consumption, but flow meters are expensive, and also may require plumbing expertise to install on certain pipes, depending on the type of sensor used; also, not all pipes providing water to a fixture may be easily exposed for installing individual flow meters. We make the key observation that water utilities are increasingly installing **smart water meters** that provide real-time access to household water consumption, and 31 million smart water meters are expected to be installed by 2016 [18]. Clearly, the already available smart water meter sensors are valuable in eliminating individual, direct sensors on each water fixture. However, these meters are installed at the water mains and only provide *aggregate whole-house water usage*, primarily for billing purposes. *Disaggregation* is the problem of deciding how much of that usage is due to individual fixtures in the home, and when these individual fixtures are being used. In this chapter, we focus on how to eliminate individual sensors on water fixtures in the home by solving the disaggregation problem on the whole house water flow data collected at the smart water meter.

There is significant existing work on the disaggregation problem. For high flow appliances such as washing machines, dishwashers, and sprinklers, existing flow trace techniques can be used on the aggregate flow data from smart meters to uniquely identify their usage time and water consumption [63]. However, monitoring the usage of multiple sinks, toilets, showers, and other fixtures that produce similar rates of flow in a home setting, is challenging. For this reason, existing disaggregation techniques require additional sensing on the water piping infrastructure, and/or a manual characterization of each water fixture [64, 55, 56, 14]. Similar to installing contact reed switches or water flow meters on each individual fixture, these techniques can be expensive, difficult to deploy, and time consuming. Also, fixtures such as multiple sinks, toilets, and showers, are most

used on an everyday basis; therefore, monitoring them is important for sustainability applications, and also for activity recognition applications.

In this chapter, we present the WaterSense system to infer fixture-level disaggregation of smart water meter data using only a simple motion sensor per room, without requiring inconvenient or expensive sensing on each water fixture. Motion sensors are inexpensive (\$5 each for X10 motion sensors [65]), and easy to install; many motion sensors can simply be placed on cabinets in an inconspicuous manner even by novice users. Also, motion sensors are already prevalent in many homes as part of home security or home automation systems. WaterSense does not require any additional sensing infrastructure on the water pipes or fixtures, and disaggregates fixtures in an unsupervised manner that does not require the collection of training data.

The WaterSense technique is based on two basic insights: 1) fixtures with similar flow signatures (e.g. identical toilets) are often in different rooms, and 2) fixtures in the same room often have different flow signatures (e.g. a toilet vs. bathroom sink). Based on these insights, WaterSense clusters all water usage events based on both flow signatures *and* motion sensor signatures, and each of these clusters represents a unique water fixture in the home. One limitation of this technique is that it is not likely to differentiate two identical fixtures in the same room, such as double sinks in a bathroom. However, such distinctions may also be less important for the purpose of water conservation decisions or activity recognition applications.

We use a three tier inference approach that first identifies water flow events from the whole house flow data, assigns each flow event to a room, and then identifies the fixture type associated with each event, to uniquely identify the usage of water fixtures in the home. A key challenge that we face in WaterSense is that individual fixtures are represented using noisy flow signatures from the water flow data, unlike the direct sensors used in AUTOLABEL; therefore, we design a Bayesian data fusion algorithm in Tier II that clusters each water flow event to a room, by combining both the motion sensor signatures and water flow signatures. Our data fusion algorithm is able to uniquely identify fixture use even in the presence of simultaneous motion activity, or lost motion sensor packets. We deployed the WaterSense system in two different, natural home environments for 7 days each and found that the system can disaggregate flow at the water mains to individual fixtures with an average 86% classification accuracy.

The rest of this chapter is organized as follows. In section 4.1, we discuss existing approaches to reduce sensor installation effort for water fixture monitoring in the home. We then present the design of the WaterSense system in section 4.2, and present its evaluation in two multi-resident homes in section 4.3.




Figure 4.1: Three tier inference algorithm used by WaterSense to find 1) water flow events 2) clusters of flow volumes that often co-occur with certain motion sensors, and 3) different types of fixtures in the same room.

4.1 State of the art

Several approaches have been proposed in the literature to eliminate or reduce sensor installation effort for monitoring water fixtures in homes. The most basic approach is to use flow signatures from the whole house flow meter data, such as flow rate, flow duration or, in the case of high flow appliances such as washing machines and sprinklers, patterns of flow to identify types of fixtures and appliances [63]. However, flow signatures alone cannot disambiguate between different instances of identical sinks, toilets, or showers in the same home. Fogarty et al. [64] use microphones installed in a basement to classify water fixture use, and achieve good accuracy in identifying high consumption appliances, but low accuracy in differentiating between different instances of the same fixture category.

Another recent approach [14] uses relatively easy to install vibration sensors on pipes to disaggregate total flow measured at a central location in a novel, unsupervised manner. Both of these techniques, however, require additional sensing of the water infrastructure, which may require access to pipes in crawl spaces or walls. Furthermore, microphones and accelerometers are more power intensive than motion sensors and would either have a short battery life or would require wired power.

Froehlich et al. [55] avoid extensive sensing by using a single water pressure sensor that samples at 500Hz, plugged into a free spigot or water outlet in the home. However, this approach requires significant training data on the order of several days in a real world setting [56] to achieve high accuracy in inferring individual water fixture events. Our WaterSense approach is an unsupervised technique that does not require training data or configuration effort from the end user.

4.2 WaterSense System Design

In this section, we describe both the underlying sensing components and the three tier inference approach used by WaterSense to infer fixture-level water usage from the whole house water flow data stream and occupancy data streams. The WaterSense algorithm uses a three tier approach, as seen in figure 4.1, and the sensor components shown in figure 4.2. In **Tier I**, we perform edge detection on the water flow data stream, matching rising edges and falling edges to compute a sequence of water flow events. In **Tier II**, we present an algorithm that groups flow events based on the motion sensors with which they often co-occur. A key challenge is that motion sensor data is very noisy due to (i) false positives from multiple residents moving in different rooms simultaneously, and (ii) false negatives from low sensitivity on the PIR sensors or wireless packet loss. To address these challenges, we design a Bayesian approach that we call **ELoc** that uses both *instantaneous* and *historical* co-occurrence of a each flow volume with each motion sensor. Finally, in **Tier III**, we differentiate the fixtures in each room into different types based on flow signatures.

Throughout this section, we will explain each of the system components in terms of an example involving two simultaneous toilet flush events. The data trace from the events is shown in Figure 4.3: the top three graphs show motion sensor data from the kitchen and two bathrooms, and the bottom graph shows the total flow levels that must be disaggregated. In this example, Tier I detects four edges with flow rates of 0.3 and 0.6 kl/hour, respectively, and identifies them as two different flow

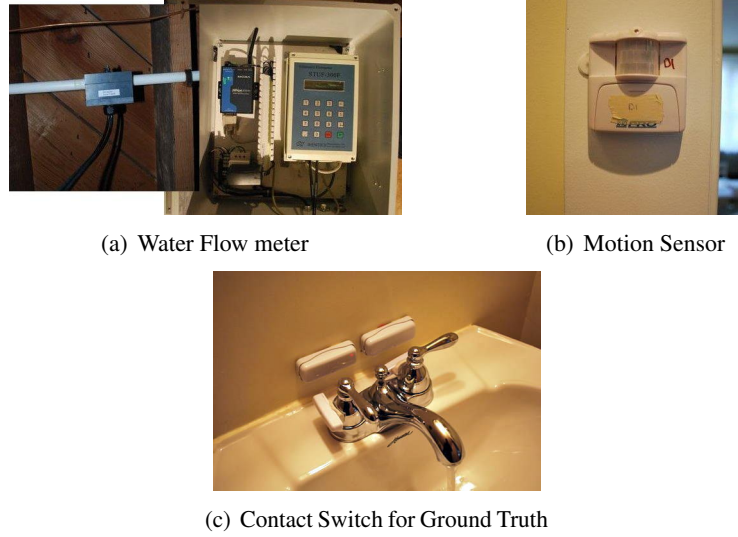


Figure 4.2: Sensors used in WaterSense and its evaluation. WaterSense uses a single flow sensor at the water mains (a) and motion sensors (b). Sensors were placed on fixtures for evaluation purposes only (c).

events. Tier II assigns the two water flow events to two different bathrooms based on their temporal proximity to motion sensor data in the two rooms, as well as the historical correlation between these flow rates and these motion sensors. Tier III identifies both flow events as toilets using the flow volume and duration.

4.2.1 Physical Sensor Components

Our WaterSense approach requires a water flow meter at the whole house water input line, and also motion sensors in each room that contains water fixtures. In our home deployments, we use a Shenitech Ultrasonic water flow meter [66] that uses the Doppler effect to measure the velocity and resulting flow of water through the pipeline. The flow meter reports instantaneous water flow (in cubic meters per hour) at a frequency of 2Hz using the home's Wi-Fi connection to transmit data. We expect that utility water flow meters being deployed in a large scale [18] in homes will have a similar setup. Figure 4.2(a) shows the installation of the flow meter in one of our home deployments. In addition to the flow meter, WaterSense requires at least one motion sensor in each room containing water fixtures. In our deployments, we use off the shelf X10 motion sensors [65] inside rooms to detect occupancy, as seen in figure 4.2(b). The X10 motion sensors send a binary ON message whenever motion is seen with a minimal damping interval of 7 seconds between ON messages.




Figure 4.3: A 6 minute data trace from Home 1 of the WaterSense deployments that shows the water flow data and binary occupancy data from two simultaneous flush events. Tier I detects four edges with flow rates of 0.3 and 0.6 kl/hour, respectively, and identifies them as two different flow events. Tier II assigns the two water flow events to two different bathrooms based on their temporal proximity to motion sensor data in the two rooms, as well as the historical correlation between these flow volumes and motion sensors. Tier III identifies both flow events as toilets using the flow volume and duration.

4.2.2 Tier I: Detecting Water Flow Events

The goal of **Tier I** is to transform the raw water flow time series $F(t)$ (in kl/hour) from our flow monitor to a sequence of timestamped *water flow events* W containing events of the form $W_i = (st_i, end_i, f_i)$, where st_i and end_i denote the start and end timestamps of the event, while f_i denotes the mean flow rate during event W_i . Tier I uses the Canny edge detection algorithm [67] on time series $F(t)$ to compute a sequence of timestamps E^t when potential edges are present. The *edge value* E_i^v for a given edge timestamp E_i^t is then computed as the difference between the median flow rate of the time intervals after and before the edge timestamp given by equation 1 below.

$$E_i^v = |\text{median}(F(E_i^t : E_{i+1}^t)) - \text{median}(F(E_{i-1}^t : E_i^t))| \quad (4.1)$$

The rising and falling edges from E are partitioned into two separate sets RE and FE respectively, matched using a min cost bipartite matching approach the Hungarian algorithm [68]). We set

each edge value to its absolute value when we partition the edges to rising and falling edges. The edge weight between a rising edge RE_i and a falling edge FE_j is defined to be $-\log(p_{ij})$, where p_{ij} is the match probability, obtained assuming a normal distribution on the relative difference between rising and falling edges:

$$p_{ij} = \mathcal{N}((RE_i^v - FE_j^v)/\min(RE_i^v, FE_j^v), 0, \sigma) \quad (4.2)$$

Additionally, we ensure that rising edges occur before falling edges in each match, and that edges from two different water flow events are not matched with each other. Specifically, we set p_{ij} to 0 if (i) $RE_i^t > FE_j^t$, or (ii) $\min(F(RE_i^t : FE_j^t)) \leq \beta RE_i^v$. Currently, we set parameters $\sigma = 0.3, \beta = 0.5$ that works well across multiple homes. Any unmatched sink edges whose corresponding ON or OFF edges coincide with a flush event, are mapped accordingly to a feasible higher flow event of opposite polarity, if such an event is detected. From the resulting matches (RE, FE) , we compute water flow events $W_i = (st_i, end_i, f_i)$ by setting $st_i = RE_i^t, end_i = FE_j^t$, and $f_i = \text{mean}(RE_i^v, FE_j^v)$. Durations T_i for each water flow events W_i are also computed as $end_i - st_i$ for use in the subsequent inference tiers.

In our example data trace in Figure 4.3, Tier I infers two distinct water flow events for the two toileting events in bathroom 1 and bathroom 2, with different start and end times. It does this by detecting two rising and two falling edges each with flow rates of 0.62 kl/hour and 0.32 kl/hour respectively.

4.2.3 Tier II: Creating Room Clusters

The goal of **Tier II** is to assign a room $\hat{r}_i \in M$ to each water flow event W_i computed in Tier I, using both W_i and the set M of relevant motion sensor data streams as input; we design the ELoc Bayesnet shown in figure 4.4 to make this room assignment. In our current approach, we only consider motion sensors in rooms containing water fixtures. Going back to our example data trace in Figure 4.3, Tier II correctly assigns water flow events W_1 and W_2 inferred by Tier I to bathroom 1 and bathroom 2 respectively. However, this assignment is not intuitive just from figure 4.3 alone; both bathroom 1 and bathroom 2 are occupied during the two flush events. Tier II first clusters all water flow events based on their duration and flow rate value, and then computes how likely these clusters are to co-occur

with the M motion sensor data streams. Event W_1 with a duration of 54 seconds and flow rate of 0.62 kl/hour belongs to a (0.6 kl/hour, 40 second) cluster that is more likely to be in bathroom 1 than bathroom 2; similarly, the other 0.32 kl/hour flush event is more likely to co-occur with the bathroom2 based on long term evidence. The ELoc Bayesnet, which uses both the instantaneous occupancy evidence from motion sensors, and also the historical room to flow signature likelihoods, accurately assigns the flush events to the two bathrooms. On the other hand, when multiple water fixtures such as flushes across multiple rooms do have closer flow signatures, such as the two 0.5-0.7 kl/hour flushes in Home 2, the Bayesnet can still use the immediate occupancy evidence to match motion sensors to flow events with high accuracy.

In the Bayesnet shown in figure 4.4, r_i denotes the room of water flow event W_i . D_i encapsulates the temporal distance between event W_i and the $|M|$ motion sensors, while $c_i \in C$ is a hidden variable that denotes the flow signature cluster to which the current event W_i belongs. Each motion sensor stream M^s consists of a sequence of timestamps when motion sensor s transmitted an ON event to our base station. We compute room assignments \hat{r}_i for each event using a maximum likelihood approach:

$$\hat{r}_i = \arg \max_{r_i} \sum_{c_i} P(r_i|c_i).P(W_i|c_i).P(c_i).P(D_i|r_i) \quad (4.3)$$

We first describe how the variables D_i and the set of water flow clusters C are computed, in order to obtain the three conditional probabilities required in equation 3 above. D_i is composed of two $|M|$ dimensional *binary vectors* D_{i1} and D_{i2} , which denote if the minimal temporal distance between the start time and end time of each water flow event and motion sensor stream M^s is below a threshold time dT (set to 12 seconds in our current approach):

$$D_{i1}^s = (\min(|M^s - st_i|) < dT) \forall s \in M \quad (4.4)$$

$$D_{i2}^s = (\min(|M^s - end_i|) < dT) \forall s \in M \quad (4.5)$$

We assume independence between the binary vectors D_{i1} and D_{i2} , and also assume that the room assignment to motion sensor j is independent of binary vectors D_i^s , when $s \neq j$. Applying

Bayes rule and our independence assumptions, we get:

$$P(D_i | r_i = j) = \prod_{k=1}^2 (P(r_i = j | D_{ik}^j) * P(r_i = j) | P(D_{ik}^j)) \quad (4.6)$$




Figure 4.4: ELoc Bayesnet used in WaterSense to cluster water flow events based on both flow volume and co-occurrence with motion sensors.

Currently, we fix the parameter $P(r_i = j | D_{ik}^j = 1) = 0.9 \forall j, k, i$, which works well across multiple homes. We currently discount $P(r_i = j)$ as a constant and set the prior probability $P(D_{ik}^j = 1)$ to be proportional to the number of ON events from each motion sensor:

$$P(D_{ik}^j = 1) = |M^j| / (\sum_{s \in M} |M^s|) \forall k, j, i \quad (4.7)$$

We compute the set C of flow signature clusters by clustering the water flow event W_i based on two features: the event duration T_i , and the event value f_i . We use quality threshold (QT) clustering with a fixed relative distance width of 0.25 on both duration and edge value [69]. QT clustering makes a *hard cluster assignment* of $q_i \in C$ to each event W_i . Thus $P(c_i | W_i) = 1$ when $c_i = q_i$ and zero otherwise. In the future, it would be interesting to explore alternative clustering algorithms with soft cluster assignments for use in the Bayesnet.

Finally, $P(r_i | c_i)$ is obtained by first assigning events W_i to rooms rd_i using D_i alone ($rd_i = \arg \max_{r_i} P(r_i | D_i)$), and then computing $P(r_i | c_i)$ using a frequency count on how many events from cluster q_i are assigned to each of the motion sensors in M . If we denote the set of events W_i such that $(rd_i = j \wedge q_i = x)$ is true to be $RC_{j,x}$, and the set of events W_i such that $q_i = x$ to be CL_x , then we get the remaining probabilities required in equation (3) as:

$$P(r_i = j | c_i = x) = |RC_{j,x}| / |CL_x| \quad (4.8)$$

$$P(c_i = x) = |CL_x|/|W| \quad (4.9)$$

From equations (6), (7), (8), (9), and our hard cluster assignments for events using QT clustering, we assign maximum likelihood room labels \hat{r}_i using equation (3). In the future, it would be interesting to learn the parameters for $P(r_i|c_i)$ using a non-linear convex optimization algorithm to maximize the likelihood of observing the sensor data, given by $L = \prod_{i=1}^{|W|} P(W_i, D_i)$. Another interesting direction is to explore alternative distance vector features D_i to potentially improve accuracy further.

4.2.4 Tier III: Differentiating Fixture Types

The goal of **Tier III** is to infer fixture types, such as a sink or a flush, for each water flow event W_i from Tier I. In the current simplistic implementation, we first infer a flush event cluster FL_j in each room j , if any, by looking for an event cluster contained in that room with an average flow rate greater than 0.3 kl/hour and an average duration greater than 30 seconds, with an event frequency of at least 10% of events from the room. Events belonging to this cluster FL_j are assigned as flush events, while the remaining low flow events are assigned as sinks. The fixture types from Tier III and the room assignments from Tier II together constitute unique water fixtures in the home, and we use these assignments to present users with feedback on the time of use of individual water fixtures distributed around a home. We also provide users with feedback on total water consumption of individual fixtures by multiplying the fixture use durations T_i of each unique water fixture with the corresponding flow values f_i .

4.3 Evaluation

4.3.1 Experimental setup

To evaluate our WaterSense approach, we deploy our system in two homes for 7 days each. Both homes had multiple residents, multiple bathrooms, and a wide array of water fixtures and appliances. Details of the two deployments are summarized in Figure 4.5. To execute WaterSense in these homes, we deployed a single water flow sensor on the water mains and a motion sensor in each room. In Home 2, one of the motion sensors in a bathroom malfunctioned during our week long deployment,

Home #	Number of Residents	Number of rooms with fixtures	Number of monitored fixtures	Pipe material and width
Home 1	2	3	5	$\frac{3}{4}$ " Copper
Home 2	4	3	5	$\frac{3}{4}$ " PEX

Figure 4.5: Home deployment details for WaterSense. Our deployments involved two homes with multiple residents. Both homes had a kitchen sink and two bathrooms with both sink and toilet.

so we used a motion sensor in an adjacent bedroom with a partial view of the bathroom in our analysis.

To evaluate the system, we deployed ZWave contact reed switches [70], as shown in Figure 4.2(c) to record the actual times that each fixture was used. We compute ground truth water consumption for each fixture by integrating total water flow into the house when each fixture was used. In the case of simultaneous toilet and sink events when we do not observe explicit sink edges, we ignore the short duration usage of the sink. These cases constitute a small fraction of events and we do not believe that it significantly changes the overall water consumption profiles. In this experiment, we installed ground truth sensors on only 3 sinks and 2 toilets in each home and only evaluate accuracy for these fixtures. Other fixtures and appliances were also used during the experiment period, and we observe several anecdotal instances of fixtures such as the shower or the sprinkler system being inferred in these homes. However, we currently limit our evaluation only to those ground truth events observed from the ZWave system. Ground truth events that do not change the total water consumption of the house are ignored, since they do not affect the water consumption feedback to the end user. In the future, it would be interesting to expand the evaluation evaluation to a larger set of fixtures and appliances.

4.3.2 Fixture monitoring accuracy

Figure 4.6 shows the confusion matrix for classifying individual fixture usage events. We observe that classification accuracy is high for the majority of fixtures in the home. Most misclassifications

Fixtures	K S	B1 S	B1 F	B2 S	B2 F
K S	8	0	0	8	1
B1 S	1	19	0	2	0
B1 F	0	1	49	0	0
B2 S	3	1	0	16	2
B2 F	0	0	0	1	14

(a) Home 1


Fixtures	K S	B1 S	B1 F	B2 S	B2 F
K S	81	10	0	3	0
B1 S	7	78	0	5	0
B1 F	1	0	85	0	5
B2 S	0	1	0	6	0
B2 F	0	0	2	1	13

(b) Home 2


Figure 4.6: Confusion matrix of WaterSense when classifying water flow events across the two homes. Confusion between fixtures of the same type in different rooms is common due to overlapping occupancy. Confusion between fixtures of different types, such as a sink and a flush, due to overlapping flow signatures, is less common. (B stands for bathroom, K for kitchen, S for sink, and F for flush)

that did occur were due to simultaneous occupancy in multiple rooms from fixtures with similar flow signatures, such as high confusion between sink usage in the kitchen and bathroom2 in Home 1. In Home 1, we find zero confusion between the flush fixtures in the two bathrooms even in the presence of simultaneous occupancy as seen in the example data trace in Figure 4.3, while in Home 2, there is about 7% confusion between the two flush fixtures. The reason is that in Home 1, the two flush fixtures have a larger difference in flow signatures (0.3 and 0.6 kl/hour flow rates), while in Home 2, the two flush fixtures have similar flow signatures of 0.6-0.7 kl/hour; the Bayesnet in Home 2 relies more heavily on the noisy instantaneous occupancy than in Home 1, resulting in slightly more errors. In general, these misclassifications do not cause significant degradation in water consumption accuracy because they are infrequent and roughly symmetric across the diagonal, and so they often cancel each other out. The average classification accuracy of 86-90% across the two homes indicates that our WaterSense approach is a promising technique to eliminate individual sensors on water fixtures for activity recognition purposes.

The true and estimated *water consumption levels* for each fixture are shown in Figure 4.7. The order of the fixtures is sorted based on their actual flow. WaterSense estimated water flow at each fixture with a median accuracy of 81.5% in Home 1 and 89.9% in Home 2. More importantly, however, WaterSense preserves the relative ordering of fixtures in terms of maximum consumption; it correctly indicates, for example, which sink or toilet causes the most water usage. End users can use the output of WaterSense to understand what the high consumption fixtures in their home are. If combined with intuitive displays that show users *when* each fixture was being used, and in which daily *activity*, we expect to provide users with actionable recommendations to save water, such as leaving the sink closed while washing or brushing, or providing low flow replacements for certain showers and flushes. For example, in Homes 1 and 2, residents might be able to infer that it might be most efficient to replace the high flow toilet in bathroom 1, while in Home 2, residents can also infer that replacing the high flow toilet in bathroom2 would have about the same effect as using the sinks in bathroom1 and the kitchen more efficiently with a better understanding of how they are used. Thus, WaterSense also shows significant promise as a cheap and unsupervised technique for monitoring the water consumption of individual fixtures in the home.



(a) Home 1



(b) Home 2

Figure 4.7: Water consumption feedback provided by WaterSense for the various fixtures monitored in the two homes, along with closely matching ground truth consumption. End users can use such a display to start considering cost effective ways to conserve water. (B stands for bathroom, K for kitchen, S for sink, and F for flush)

4.4 Limitations

A key limitation of our WaterSense evaluation is that it does not consider water flow events from high flow appliances such as the washing machine, dishwasher, and the shower that are not instrumented to measure ground truth usage. In the future, it would be interesting to perform a complete evaluation that includes all water fixtures and appliances, to address potential challenges posed due to simultaneous use of high and low flow appliances. As more high flow appliances are considered in the evaluation, Tier III of the WaterSense algorithm needs to be improved to handle complex appliances such as washing machines and dishwashers, using insights from the flow profiling literature [63].

Another limitation of our evaluation is that our ground truth measurement system currently combines contact reed switches with the whole house water flow data to measure the water consumption of individual fixtures; such a system is prone to errors when water fixtures are being used simultaneously, since the water flow increase is not proportional to the number of fixtures as more fixtures are used concurrently. Also, to address false positives in the reed switch sensors, we ignore reed switch events that do not cause an increase in the whole house water flow; such an approach could classify certain true positive reed switch events as false positives, if they are low flow events that are not detected at the whole house flow meter. Conversely, flow events that do not cause a corresponding firing in any of the contact reed switches are assumed to be from uninstrumented high flow appliances such as the washing machine or the dishwasher; due to false negatives in the reed switch sensors, such an approach may cause certain true positive events from instrumented fixtures such as the sink and the flush to be ignored as being from an uninstrumented fixture. To handle these challenging problems with ground truth measurement, a comprehensive, but more expensive system consisting of a flow meter at every water fixture could be used in the future to provide accurate ground truth.

Finally, in our evaluation of WaterSense, we only consider data from two homes over a short term period of 7 days each, with two bathrooms and a kitchen containing water fixtures in each home. In the future, it would be important to extend the evaluation to include more homes and apartments, and also perform the evaluation over a longer term. It would be interesting to evaluate if WaterSense achieves high accuracy in such a larger scale evaluation, or if lower accuracy is achieved in certain homes due to similar fixture types and higher simultaneous use.

4.5 Conclusions

In this chapter, we presented the unsupervised WaterSense approach to infer the fine-grained fixture level breakdown of water consumption in homes by effectively combining cheap occupancy sensors and whole house water flow meter data. Our WaterSense approach is evaluated in two homes for 7 days each, and is able to accurately infer both the time of use and water consumption of individual fixtures.

Our three tier inference approach for WaterSense may be used to provide the underlying water fixture usage input to the three tier AUTOLABEL inference algorithm discussed in the previous chapter. In AUTOLABEL, each water fixture is sensed by a unique binary sensor. On the other hand, in WaterSense, we first need to detect water flow events from the whole house flow data using an event detection tier, since there is no longer a binary sensor on each water fixture. For clustering sensors to rooms, AUTOLABEL leverages a unique temporal distance computed between every pair of distinct sensors. In contrast, we need to use a Bayesian clustering approach in WaterSense to effectively fuse the noisy flow signature and motion signature evidence, to cluster flow events from the water flow data to rooms containing motion sensors. Also, the flow signatures in WaterSense are able to identify the actual fixture type, such as a sink or a flush. On the other hand, AUTOLABEL only relies on temporal features of usage to determine the fixture type in tier III. In the future, the fixture type information from WaterSense can be leveraged in the AUTOLABEL algorithm to determine the room label in the room classification tier (Tier II of the AUTOLABEL algorithm).

In our current implementation, we only use off the shelf binary PIR sensors in each room containing a water fixture. In the future, it would be interesting to consider non-binary PIR sensors that report the actual intensity of infrared changes, as opposed to the existing binary motion sensors; with non-binary PIR sensors, it would be possible to perform a joint clustering of the flow rate and the infrared intensity, to achieve higher accuracy in identifying individual fixtures within a room, and also better disambiguate cases where there is simultaneous occupancy in multiple rooms. It would also be interesting to explore how weak biometric sensors used for resident identification in rooms can be used along with the occupancy sensors in the WaterSense approach to better disambiguate water fixture usage.

Chapter 5

Reducing sensor installation effort to monitor electrical fixture usage

In chapter 4, we presented the WaterSense solution to eliminate direct sensors on individual water fixtures in the home. In addition to water fixtures, the usage of individual electrical fixtures in the home is a key input for recognizing resident activities. For example, the use of appliances such as the stove and microwave indicate cooking, while the use of appliances such as the dishwasher or the washing machine indicate hygiene activities [28, 27, 26, 37]. In our evaluation of our AUTOLABEL activity recognition approach in chapter 3, we used magnetic contact reed switches on individual electrical fixtures in the home. As discussed in chapter 4, reed switches can be time consuming for a novice to install, and also inconvenient due to the presence of conspicuous sensors on electrical appliances used every day. Also, reed switches do not indicate how much energy is being consumed by the light fixture under consideration. A naive alternative to reed switches is to install an in-line current sensor at each electrical switch; however, based on our own experiences, such an effort consumes more than 20 man-hours on average per home to install, and also inconveniences the end user by requiring power to be shut off while the installation is ongoing. In this chapter, we present new approaches for eliminating direct sensors on individual electrical fixtures in the home.

Similar to our WaterSense approach, it is possible to leverage the smart meters being deployed in homes by utility companies to monitor the use of individual electrical fixtures. For example, non-intrusive load monitoring (NILM) approaches [71] have long been used to identify the usage of

major appliances by analyzing their electrical power signature, using only a single power meter on the electrical mains of the home; this approach achieves high accuracy for high power appliances such as the stove, the washing machine, or the microwave. However, NILM is not effective for low-power appliances (30-100W), since several low power appliances have similar wattages in 30-100W range, and the wattage of a low power appliance is similar to the electrical noise levels of a typical home. In this chapter, we choose to focus on the highest power consumers among low power appliances, namely light fixtures.

Light fixtures play an essential role in the home, and monitoring their usage can serve activity recognition systems targeted at both sustainability and medical applications. For example, task lighting near a chair, countertop, or table can be a good indicator of medically relevant activities of daily living (ADLs) such as reading, preparing food, and eating dinner. In fact, light fixture usage may be the primary indicator of certain activities, at least from a sensing perspective. For example, the status of the light in a bedroom may be an indicator of either reading or sleeping [27].

Monitoring light usage is also important from an energy conservation viewpoint. Light fixtures account for 11% of energy consumption in US homes on average [72] and 25% in California because heating and cooling energy is lower due to the mild climate [73]. This energy usage is equivalent to that used by 15 million homes or about 20% of all automobiles on the road in the US. Homeowners can reduce lighting energy by installing energy-efficient CFL or LED bulbs, daylight harvesting hardware, or occupancy-automated lighting controls. For example, when both occupancy and light sensors are available, researchers have shown that smarter control of light fixtures can yield savings of 50-70% through light fixtures fitted with actuators. However, these solutions are costly relative to the energy saved: some of the homes we analyzed contained over 50 bulbs, including specialized lighting (e.g. recessed lighting or under-cabinet lighting) that is particularly costly to retrofit. Homeowners need concrete data about individual light fixture usage to quantitatively assess the expected return on investment of any given solution. Also, past studies have shown that up to 15% energy savings can be achieved for individual appliances, if residents have a precise understanding of their energy costs [19, 20, 21].

Unfortunately, light fixture monitoring remains a challenge today, since approaches such as NILM are not effective in the wattage range of a typical incandescent light bulb(30-100W). Recently, several alternatives to NILM have been developed, but each has important limitations. For example,

they require a sensor to be installed on every fixture [11, 13, 12, 74], require a manual training and labeling period [58], or are limited to light fixtures that involve switched-mode power supplies (SMPS) [57] such as those with CFLs or dimmer switches.

A key challenge with monitoring light fixtures using the power meter is the high level of noise; using the power meter alone, we observed that only 15% of the low power events corresponded to light fixture usage in the home. Similar to the WaterSense approach, we first explored combining the motion sensor data with the power meter data to eliminate noisy events from the power meter, and disambiguate light fixtures in different rooms with the same wattage. However, we found that such an approach still produced a significant number of false positives when trying to identify genuine light fixture events. The high noise in the power meter data was not effectively filtered by motion activity, since motion activity durations were very long, and only sometimes correlated with light fixtures turning ON and OFF. To effectively reduce the number of false positives, we explored using alternative sensors to better eliminate false positives from the power meter data, compared to motion sensors.

In this chapter, we present a new approach called LightSense that uses data fusion between the whole house power meter and a light sensor in each room, to cheaply, easily, and accurately infer light fixture usage. LightSense effectively captures both the electrical signature and the light signature of each fixture in an unsupervised manner. For a typical home, this approach requires only 3-8 more sensors than conventional NILM techniques. Furthermore, light sensors are inexpensive (\$ 0.5-1) and can be easily installed without touching the electrical wiring. Also, these light sensors can be piggybacked on the motion sensor devices in each room used by our WaterSense approach.

However, light sensors alone are not sufficient to detect fixture usage: the light levels produced by a typical fixture are similar to or even smaller than the light levels caused by window shades opening and closing, people walking past the sensor, and clouds passing overhead. Light fixtures account for only a small fraction (25%) of the light level changes in a room. To address this problem, LightSense exploits several insights about light fixture usage: 1) noise on the power lines and noise in light are typically independent 2) light fixtures typically exhibit a similar power/light signature over time, even when using dimmers 3) light usage always occurs in ON/OFF pairs: a light cannot be turned ON twice in a row. In this chapter, we present a probabilistic data fusion framework that

combines power and light data streams, automatically learns the long-term signature of each light fixture, and accurately recognizes the real-time usage events.

LightSense is completely unsupervised: it does not require a training process in which users manually label light fixture usage. Furthermore, it does not require user configuration: users do not need to indicate the number of fixtures in a home or their individual wattages. These parameters are automatically inferred.

We envision two types of LightSense deployments. In a short-term *energy audit*, light sensors would be placed on a shelf or countertop in each room for a period of several days or weeks to assess the total energy usage of each light fixture. Power data would be collected from a smart power meter, which are already deployment in many parts of the US by power utility companies for billing purposes. For longer-term activity recognition applications, the light sensors would simply be integrated with other sensors. For example, several commercial motion sensors already have integrated light sensors for calibration purposes [65, 70]. If a motion sensor is placed in every room for activity recognition purposes anyway, the light sensor data can be used by LightSense to infer light fixture usage.

We evaluate LightSense with an *in-situ* deployment of power meters and light sensors in four multi-resident homes over 10 days each, monitoring a total of 41 light fixtures across 32 rooms. We deployed wireless light switches and smart plugs to collect ground truth information about light fixture usage. We found that the power meter alone detects light fixture events with only 15% precision, and the light sensors alone achieve only 25% precision. By combining the two data streams, however, the LightSense data fusion algorithm detects ON/OFF events with 81% recall and 86% precision. We also demonstrate that LightSense can disaggregate whole-house energy usage into fixture-level energy usage, estimating the total energy usage of most the commonly used light fixtures with 91.1% accuracy on average.

The rest of this chapter is organized as follows. In section 5.1, we discuss the state of the art in cheaply monitoring light fixture usage in homes. In section 5.2, we discuss the LightSense data fusion algorithm. In section 5.3, we discuss the experimental setup and home sensor deployments used to evaluate LightSense. In section 5.4, we discuss the accuracy results for LightSense in monitoring individual light fixture energy consumption, and in section 5.5, we present the accuracy of LightSense in identifying individual light fixture ON and OFF events for activity recognition purposes, and

Technique/ Feature	Easy to install?	Cost	Energy feedback?	Accuracy	User training?	Applicability?	In-site evaluation?
Electrisense	Yes	High	No	High	No	CFLs	No
Using mechanical properties of switches	Yes	High	No	High	Yes	Incandescent	No
NIALM	Yes	Low	Yes	Very Low	Yes	All	Yes
Direct sensing	No	High	Yes	Very High	No	All	Yes
Ambient sensing	No	High	Yes	High	No	All	No
Light Sense	Yes	Low	Yes	High	No	All	Yes

Figure 5.1: Comparison of existing techniques to infer usage and power consumption of light fixtures

provide an explanatory analysis to show why LightSense achieves high accuracy in spite of noise from the light and power data streams.

5.1 State of the art

Several approaches have been developed for monitoring light fixture usage, but each has important limitations. Figure 5.1 lists the main techniques in this area, and compares them to LightSense in terms of several important metrics.

Non-Intrusive Appliance Load Monitoring (NIALM) [71] techniques can be used to identify the usage of electrical fixtures in the home based on electrical signatures. These signatures can be extracted using only a single power meter [17], which is already available in many homes today [16]. Some appliances have a unique profile of real and reactive power, while other appliances such as washing machines and dishwashers exhibit characteristic electrical patterns over time. However, NIALM techniques not effective for light fixtures that exhibit constant, low power values due to the large number of similar, low power appliances in homes, and due to low power state transitions from complex appliances such as the television or the HVAC system [71].

An alternative to NIALM is use *in-line current sensors* to directly measure the amount of energy consumed by individual appliances. Examples include pluggable power meters [70, 11] to monitor lamp fixtures, and smart switches [70] to monitor lights controlled by wall switches. However, direct sensing a large number of sensors to be integrated with the home’s electrical wiring, which is costly in terms of both hardware and installation time. In our own deployments across four homes, for example, installing in-line current sensors for ground truth consumed approximately 19 man hours

per home on average, while the LightSense installations only consumed an average of 25 minutes per home.

Recently, new solutions have been developed that require only a single sensor on the power line. Gupta et al propose the *Electrisense system* [57], which uses an easy to install, plug-in sensor that leverages unique high frequency EMI(Electromagnetic Interference) signals on the power line to identify light fixtures. However, this approach is limited to appliances that use switched-mode power supplies (SMPS) such as lights with CFL bulbs or dimmer switches; it does not work for conventional incandescent bulbs. Patel et al [58] propose a related approach to identify *mechanically switched appliances* in homes based on a unique high frequency voltage transients. However, this approach requires users to manually train the system by labeling ON and OFF events of individual appliances. If lights are added or moved throughout the house, training must be performed again. Furthermore, the authors only evaluate the ability to classify manually labeled events that are explicitly input to the classifier(recall and classification accuracy). It is unclear if the reported upper bound classification accuracy for these approaches of 75-95% will be retained in a realistic, in-situ evaluation, where *false positives* in the power line signals could reduce precision. Moreover, neither of these approaches can currently infer the energy consumption of individual fixtures; they can only identify which fixture was used. LightSense is a completely unsupervised technique that can identify fixtures and measure their energy consumption. LightSense is also evaluated in a *in-site* experiment.

The systems most similar to LightSense include *ambient sensing* approaches that fuse data from whole-house power monitors with environmental sensors deployed throughout the house. For example, Viridiscop [12] uses one specialized sensor node per appliance containing light, acoustic, and magnetic sensors with whole house power meter data to infer the energy costs of individual appliances in the home. However, the Viridiscop algorithm makes the strong assumption that the power consumed by unmonitored appliances is constant, essentially requiring complete coverage of every electrical fixture and appliance in the home. A contactless sensing system proposed by Rowe et al [13] also requires specialized environmental sensors per appliance, and additionally uses high frequency circuit level power meter data, which requires high installation effort at the breaker box. Both of these approaches require far more sensors and a more involved installation process than LightSense. LightSense can detect light fixture usage without the need to monitor any other fixtures or appliances in the home. Further, it does not require a sensor on each fixture, but only requires full

coverage of all light fixtures of interest (typically, one sensor per room). More recently, Delaney et al proposed a system to save energy from lighting based on smart control in an office environment, and presented savings of 50-70% [22]. However, they used direct sensors and actuators on each light fixture, and do not address the problem of effectively inferring the energy usage of each light fixture.

5.2 LightSense Design

The LightSense algorithm uses a three-tier data fusion framework to combine raw light sensor and power meter data and generate the usage times and energy costs of individual light fixtures. The framework is illustrated in figure 5.2. The Tier I applies edge detection algorithms to identify both rising and falling edges in both data streams. LightSense uses several novel edge detection algorithms tailored to the noise typical of light and power timer series. Tier II combines the edges detected in both data streams to generate a single series of events. Tier II applies three basic principles to decide which edges to keep: a) the edge must appear in both the light and power data b) the light and power values must be typical for that room c) there must be a matching between rising edges and falling edges. Once the light fixture events are identified, Tier III uses a clustering algorithm to infer the individual light fixtures present in each room, their usage times, and their total energy costs. Figure 5.3 shows a sample data trace from 6AM to 4PM in one home. The output of Tier I produces a large number of false edges for both the light and power data, and the data fusion algorithm in Tier II eliminates almost all of them to identify only 4 correct ON/OFF events.

5.2.1 Tier I: Light and Power Edge detection

The goal of Tier I is to detect rising and falling edges in the power and light data streams, despite several types of noise. The inputs to Tier I are the time series data from each light sensor i denoted by L^i , and the time series data from the whole house power meter P . Using this input, the edge detection algorithms output a sequence of rising and falling edges for each light sensor i , denoted by RL^i and FL^i , respectively, and a sequence of rising and falling edges for whole house power consumption, denoted by RP and FP , respectively.




Figure 5.2: Three tier data fusion framework used by LightSense to combine data streams.

Power Edge Detection

To detect whole house power consumption edges from the power meter time series P , we apply a sliding window based edge detection algorithm, that computes all possible power edges with a maximum window bound. In particular, each power edge e is defined by a time window $[e_s, e_f]$, and the edge value is given by $e_v = (P[e_f] - P[e_s])$. We place two main constraints on time windows which can constitute a power edge e :

1. $(e_f - e_s) \leq \text{maxwin}P$
2. $|e_v| \geq dP$

Condition 1 ensures that all edges have a time window length bounded by $\text{maxwin}P$. A small parameter for $\text{maxwin}P$ parameter ensures that we eliminate a significant number of slow power intensity changes that do not originate from artificial lighting. In our current implementation, we fix $\text{maxwin}P = 5$ seconds across all homes. *Condition 2* determines the lowest wattage light fixtures detected by LightSense, parameterized by dP ; in section 5.5.3, we evaluate the recall precision tradeoff as dP is varied.

Our power edge detection algorithm is a multi-edge detection approach, since multiple power edges could simultaneously be detected at any time instant from the different time window lengths bounded by $maxwinP$. We adopt this approach because the power meter data typically contains several closely spaced simultaneous low power events, and a single edge window is insufficient to achieve high recall; a larger edge window is ineffective when there are closely spaced events ON or OFF events, while a smaller edge window does not capture slowly increasing power events from dimmable lights as reported by the power meter, which reports rms power every second from the AC current and voltage inputs. We address the problem of pruning power edges in section 5.2.2, by leveraging the joint light and power signatures of light fixtures. Also, we bound the maximum power edge magnitude $|e_v|$ to $1000W$ to prevent interference from high power appliances such as the stove or the washing machine. The edges output by our edge detection algorithm are partitioned to rising and falling power edges RP and FP , respectively, and the edge value e_v is set to $|e_v|$ in the partitioned subsets.

Figure 5.3 shows the timestamps of edges output by the power edge detection algorithm. We observe that the power edge detection approach shows poor precision in detecting light fixture events, due to false positives from numerous low power appliances, and low power state transitions of complex appliances such as the HVAC system or the television.

Light Edge Detection

To detect light edges from each light sensor stream, we design the light edge detection algorithm outlined below. Figure 5.3 shows the timestamps of rising (ON) and falling (OFF) edges output by the two steps in the light edge detection algorithm. The first step uses a sliding window-based edge detection algorithm similar to the power edge detection algorithm to eliminate noise due to slowly changing light events, such as daylight changes from the movement of the sun. The second step uses *three novel adaptive noise filters* to filter the typical noise events seen in home light sensors.

```

for each light sensor  $i \in S$ 
   $LS^i = medfilt(L^i, window)$ 
   $E^i = windowdetect(LS^i, maxwin, dL)$ 
   $F^i = null$ 
for each  $e \in E^i$ 

```

```

 $b = \text{noise\_filter}(e) \wedge \text{onoff\_filter}(e) \wedge \text{consensus\_filter}(e)$ 

if  $b == 1$  then add  $e$  to  $F^i$ 

partition  $F$  to  $(RL, FL)$ 

```

In the algorithm above, for every light sensor $i \in S$, where S is the set of light sensors in the home, we first perform median filtering on the light sensor time series L^i with a fixed *window* size of 10 samples across all four homes. Median filtering eliminates impulse noise from our light sensors. We then apply a window-based edge detection algorithm *windowdetect* on the smoothed time series LS^i to output a sequence of edges E^i . This algorithm uses the two conditions mentioned above for the window-based power edge detection algorithm, with new parameters dL and $maxwinL$ instead of parameters dP and $maxwinP$. Setting a small bound on the $maxwinL$ parameter ensures that we eliminate a significant number of slow light intensity changes that do not originate from artificial lighting. In our current implementation, we fix $maxwin = 2.5$ seconds across all homes, since that is the maximum time required for typical residential dimmer switches to reach their target light intensity level. Parameter dL determines the lowest intensity of light edges detected by LightSense; in section 5.5.3, we evaluate the recall precision as dL is varied. We also add two additional constraints on edges $e \in E^i$ to ensure that multiple light edges are not detected simultaneously:

1. $\forall d \in E^i, (e \neq d) \rightarrow ([e_s, e_f] \cap [d_s, d_f] = \{\})$
2. $\forall [x, y], (|L^i[y] - L^i[x]| \geq dL) \wedge ([x, y] \subset [e_s, e_f]) \rightarrow (e_v * (L^i[y] - L^i[x])) > 0$

The above conditions together ensure that the edge window sizes are adjusted automatically depending on the variation in the light signal; smaller edge windows are used when light intensity changes are dense, while larger edge windows are used when the light intensity changes are sparse. Due to space constraints, we do not discuss how the *windowdetect* algorithm ensures the constraints above to output E^i , but we note that it is a straightforward algorithm that is in $O(|L^i|)$.

For each edge $e \in E^i$ output by the *windowdetect* algorithm, we apply three additional adaptive filters to eliminate noisy edges, namely the *noise_filter*, the *onoff_filter*, and the *consensus_filter*; in figure 5.3, we see that these adaptive filters eliminate a significant number of noisy edges output by the *windowdetect* algorithm.

Firstly, lights do not turn on and off very frequently; if we detect a large number of similarly sized light edges, we assume they are ambient noise. For each edge e , the *noise_filter* algorithm first




Figure 5.3: Sample Data Trace from a bedroom light sensor and whole house power meter from 4AM to 4PM in House 2 of our LightSense deployments, along with ground truth fixture events. LightSense eliminates false positive light and power Edges by performing data fusion and matching.

counts the number of comparable, temporally adjacent edges x within time dT_a of edge time e_s , whose magnitude is at least α times the current edge value e_v under consideration. The filter returns *false* if the number of such comparable edges exceeds a *maxcount* parameter. We fix parameters $dT_a = 20$ minutes, $maxcount = 20$, and $\alpha = 0.8$ across all homes.

Secondly, when people or clouds pass by the light sensor, a rising edge and falling edge occur closely together in time; we filter these pairs of events out. For each edge e , the *onoff_filter* algorithm returns false if an *opposing, comparable, temporally adjacent* edge y exists within time window dT_b of edge time e_s , whose magnitude y_v is at least α times the current edge e under consideration, and whose polarity is *opposite* to edge e . We fix parameter $dT_b = 12$ seconds.

Thirdly, redundant light fixture events from adjacent rooms are detected in multiple light sensors. So, when we observe light edges in multiple light sensors simultaneously with the same polarity, we only retain the light edge with maximum intensity, since that edge is most likely from the sensor in the same room as the light fixture. Assuming a coarse time synchronization among the multiple light sensor streams based on timestamps at the receiver base station, the *consensus_filter* eliminates

redundant light edges within one second of each other. Finally, the filtered light edges are partitioned to rising and falling edges based on their polarity to RL and FL , respectively.

5.2.2 Tier II: Data Fusion and Matching

The goal of Tier II is to combine the low precision light and power edges from Tier I to compute a sequence of *matched* ON-OFF events M^i for each light sensor i . Each ON-OFF event $m \in M^i$ is defined by a four tuple (m_s, m_f, m_p, m_l) ; m_s and m_f denote the ON and OFF time, respectively, while m_p and m_l denote the average power consumption and light intensity increase for event m . Firstly, we use *data fusion* to eliminate temporally isolated power or light edges as false positives. Secondly, we use *Bayesian matching* to prune unmatched ON or OFF light and power edges as false positives. Finally, we use a novel *Bayesian clustering* approach to accurately assign power costs to each ON-OFF event, by automatically learning the typical power and light signatures in a room.




Figure 5.4: The lpbayes Bayesnet used by LightSense to find the maximum likelihood clusters (top nodes) given the power and light edges (lower nodes).

Data Fusion

The Data Fusion step combines the rising and falling power edges $\{RP, FP\}$ and light edges $\{RL^i, FL^i\}$ to compute a set fused light-power edges $\{RLP^i, FLP^i\}$ for each light sensor $i \in S$. We first describe how the rising light-power edges RLP^i are computed from RP and RL^i . We first compute a power edge set PS_e for each light edge $e \in RL^i$, where PS_e denotes all the power edges from RP within time dT_{add} of the light edge timestamp e_s . If no power edges are found, we discard the light edge e , otherwise we add the light-power edge (e, PS_e) to RLP^i . Thus, temporally isolated light or power edges are eliminated as false positives, while temporally adjacent light and

power edges are added to RLP^i . We similarly compute the falling light-power edges FLP^i from FP and FL^i . We currently set dT_{add} to 12 seconds across our deployments to account for time synchronization errors between the light sensor and power meter data streams. Since the power meter data is noisy, *multiple power edges in PS_e are typically associated with a single light edge e* ; the problem of assigning the correct power edge to each light edge is addressed as part of our Bayesian Clustering step below.

Bayesian Matching

For each light sensor $i \in S$, the Bayesian matching step matches *rising* light-power edges RLP^i to *falling* light-power edges FLP^i to compute the sequence of matched light fixture events M^i . Bayesian matching prunes false positive edges from RLP^i and FLP^i that cannot be matched with high probability. The algorithm is shown below:

```

for each light sensor  $i \in S$ 
   $Weights^i = null$ 
  for each edge  $(d, PS_d) \in RLP^i$ 
    for each edge  $(e, PS_e) \in FLP^i$ 
       $c1 = (d_s > e_s)$ 
       $c2 = \min(L^i[d_s : e_s]) > \gamma * d_v$ 
      if  $(c1 \vee c2)$ 
         $pmatch_{d,e}^i = 0$ 
      else
        compute  $lprob_{d,e}^i, pprob_{d,e}^i$ 
         $pmatch_{d,e}^i = lprob_{d,e}^i * pprob_{d,e}^i$ 
         $Weights^i[d, e] = -\log(pmatch_{d,e}^i)$ 
   $M^i = \text{Hungarian}(Weights^i)$  with weight threshold  $wmax$ 

```

The Bayesian Matching algorithm shown above first constructs a weighted bipartite graph with edges from the set of rising light-power edges RLP^i , to the set of falling light-power edges FLP^i . The edge weight $Weights^i[d, e]$ between any rising edge $d \in RLP^i$ and any falling edge $e \in FLP^i$ is set to $-\log(pmatch_{d,e}^i)$, where $pmatch_{d,e}^i$ represents the probability of edge d being matched to edge e . We perform a min-cost bipartite matching on the weighted bipartite graph

represented by $Weights^i$, which returns an optimal matching M^i with maximum match likelihood, i.e. $M^i = \arg \max_x (\prod_{(d,e) \in x} pmatch_{d,e}^i)$.

In the algorithm above, we see how the match probabilities $pmatch_{d,e}^i$ are computed. Firstly, we set $pmatch_{d,e}^i = 0$ if any of the two conditions $c1$ or $c2$ are satisfied. Condition $c1$ ensures that rising edges in an ON-OFF event occur *before* falling edges. Condition $c2$ ensures that rising and falling edges from two different ON-OFF event pairs are not matched together by leveraging the observed additive nature of light intensities in a room; if a +100 rising edge and a -100 falling edge are separated by a time interval where the total light intensity is only 5, then the two edges are likely from different ON-OFF events. In the algorithm, γ is set based on empirical experiments to 0.8 across all our deployments. If conditions $c1$ and $c2$ are false, the two edges d, e under consideration could potentially be matched.

As shown in the matching algorithm, the match probabilities $pmatch_{d,e}^i$ are computed as $lprob_{d,e}^i * pprob_{d,e}^i$, where the light match probability $lprob_{d,e}^i$ denotes the probability that the light edges d and e are matched, while the power match probability $pprob_{d,e}^i$ denotes the probability that the two corresponding *power edge sets* PS_d and PS_e are matched.

We compute the light match probability $lprob_{d,e}^i$ as the likelihood that ON and OFF light edges d, e belong to the same light fixture; we observe that ON and OFF light edges from the same two-state light fixture have similar edge values. To leverage this observation, we first cluster all the *light edge values* from RLP^i and FLP^i using the QT (Quality Threshold) clustering algorithm [69] with a *relative distance error threshold* of 0.25 on the light edge values. QT Clustering outputs a set of light edge clusters CL^i with a fixed cluster assignment for each edge. We assume each light edge cluster has a normally distributed light intensity distribution, and compute the mean and standard deviation for each light cluster $cl \in CL^i$ as cl_{mean} and cl_{std} , respectively, from the edges assigned to cl . The probability $lprob_{d,e}^i$ that the two light edges d, e belong to the same light cluster $cl \in CL^i$ is then computed as:

$$lprob_{d,e}^i = \sum_{cl \in CL^i} \mathcal{N}(d_v, cl_{mean}, cl_{std}) * \mathcal{N}(e_v, cl_{mean}, cl_{std}) \quad (5.1)$$

Computing the power match probabilities $pprob_{d,e}^i$ is less straightforward since there are typically multiple rising and falling edges in the power edge sets PS_d and PS_e . To simplify this problem, we choose the assignment $x \in PS_d, y \in PS_e$ that maximizes $pprob_{d,e}^i$ as follows:

$$pprob_{d,e}^i = \max_{x,y} (vprob_{x,y}^i * p^i(x, d_v) * p^i(y, e_v)) \quad (5.2)$$

In equation 2, the power value match probability $vprob_{x,y}^i$ denotes the probability that the two power edges x, y are from the same light fixture, while the joint light-power probabilities $p^i(x, d_v), p^i(y, e_v)$ denote the probability that a light fixture with light edge value d_v or e_v has a power consumption of x or y , respectively. To compute the power value match probability $vprob_{x,y}^i$, we leverage the observation that two-state light fixtures produce ON and OFF power edges with similar values. We cluster all the *power edge values in edge sets* from RLP^i and FLP^i using the QT (Quality Threshold) clustering algorithm with a *relative distance error threshold* of 0.25 to obtain the set of power edge clusters CP^i ; from CP^i , $vprob_{x,y}^i$ is computed similar to $lprob_{d,e}^i$ in equation 1. In equation 2, the light-power probability $p^i(x, d_v)$ denotes the joint probability of observing power edge value x together with light edge value d_v from light sensor i . We use the **lpbayes** Bayesian Clustering approach described below to compute the joint light-power probability values.

Bayesian Clustering

Our Bayesian Clustering approach computes the joint probability $p^i(p_k, l_k)$ of observing any power edge value p_k and any light edge value l_k together in the same light-power edge from light sensor $i \in S$. The joint probability must effectively capture the typical light and power signatures that occur together in a room. For example, consider an example ON-OFF match from a 100W light bulb, containing a rising edge of $d = (40L, [40W, 103W, 500W])$ and a corresponding falling edge of $e = (41L, [40W, 95W, 1000W])$ sensed by light sensor i . We expect $lprob_{d,e}^i$ to be high for this match, since the two light edges are very likely to be from the same 40L light cluster. Among the multiple power edge assignments possible, the highly likely assignments based on $vprob^i$ from equation 2 are $(40W, 40W)$ and $(103W, 95W)$. If we use only $vprob^i$, we incorrectly assign the $(40W, 40W)$ power edges to the ON-OFF event, since they are more likely to be from the same cluster than the $(103W, 95W)$ edges which are more noisy. However, by leveraging the joint probability of occurrence

of the light and power meter data streams using `lpybayes`, we still make the correct 100W assignment to the ON-OFF match because the 103W event has a much higher likelihood than a 40W event of co-occurring with the 40L light event, i.e. $p^i(103W, 41L) \gg p^i(40W, 41L)$.

Our Bayesian Clustering approach uses the **lpybayes** Bayesnet shown in figure 5.4 to compute the joint probabilities $p^i(p_k, l_k)$; in the Bayesnet, cl^k and cp^k are hidden random variables that denote the cluster to which the light edge value l_k and power edge value p_k belong to, respectively. The Bayesnet effectively captures the probabilistic relationship between individual light or power values and their cluster assignments, and more importantly, the conditional relationship between the light and power clusters, indicating how likely an element from power cluster cp^k occurs given that an element from light cluster cl^k has occurred. From the basic properties of Bayesnets, the joint light-power probabilities $p^i(p_k, l_k)$ are computed as follows:

$$p^i(p_k, l_k) = \sum_{cl^k} \sum_{cp^k} p(l_k | cl^k) * p(p_k | cp^k) * p(cp^k | cl^k) \quad (5.3)$$

The conditional probabilities $p(l_k | cl^k)$ and $p(p_k | cp^k)$ are computed using the Gaussian distributions computed for the light and power edge clusters during the QT clustering process described in the previous section, i.e. $(p(l_k | cl^k) = \mathcal{N}(l_k, cl_{mean}^k, cl_{std}^k))$ and $(p(p_k | cp^k) = \mathcal{N}(p_k, cp_{mean}^k, cp_{std}^k))$. The conditional probability $P(cp_k = a | cl_k = b)$ is computed by frequency counting the number of individual light edges z assigned by QT clustering to *light cluster* b that have *at least one* edge in PS_z assigned to power cluster a . In other words, if we denote the set of light-power edges assigned to light cluster b as CA_b , and the set of light-power edges in which at least one element in the power edge set is assigned to power cluster a as CA_a , then:

$$p(cp_k = a | cl_k = b) = |CA_a \cap CA_b| / |CA_b| \quad (5.4)$$

The power assignments x, y that maximize equation 2 are fixed as the optimal power assignments for the current match (d, e) under consideration. From equations 1, 2, 3, and 4, we calculate the match probabilities $pmatch_{d,e}^i$. After computing the match probabilities, the edges from RLP^i and FLP^i are matched using the min-cost bipartite matching algorithm (Hungarian algorithm [68]). We limit the maximum match weight $wmax$ to 70 in our deployments to eliminate highly noisy matches which are unlikely to be true light fixture events. For each match $m \in M^i$, the start and end times

m_s and m_f are set to the timestamps of the rising and falling edge, respectively. The light intensity value m_l and the power value m_p are set to the mean of the light and optimal power edges in the rising and falling edge of match m , respectively.

5.2.3 Tier III: Fixture Identification

The goal of Tier III is to identify the individual light fixtures in each room, their nominal wattage values, usage times and total energy costs. Using the ON-OFF matches M^i from Tier II as input, Tier III outputs a set of light fixtures LF^i for each light sensor i ; for each light fixture $lf \in LF^i$ we output: (a) the sequence of ON-OFF event times given by $lf_m \subseteq M^i$, (b) the nominal wattage value lf_w , and (c) the total energy consumption lf_e .

To compute LF^i , we cluster the ON-OFF matches M^i based on the match light intensity m_l using QT clustering with a relative distance error threshold of 0.25. For each light fixture cluster $lf \in LF^i$ output by QT clustering, we compute (i) the usage times based on the start and end times of the matches $lf_m \subseteq M^i$ assigned to cluster lf , (ii) the total energy consumption based on the ON-OFF event timings and power consumptions as $lf_e = \sum_{m \in lf} m_p * (m_f - m_s)$, and (iii) the nominal wattage value of the light fixture as the *median* of the power assignments m_p made to matches $m \in lf$. We perform a simple threshold based filtering to eliminate any degenerate light fixtures $df \in LF^i$, if the number of events assigned to the light fixture df is less than 2, or if the total energy contribution of the light fixture to the room containing the light sensor i is less than ϵ %, i.e. if $df_e < (\epsilon/100) * (\sum_{y \in LF^i} y_e)$. We set ϵ to 10% in our implementation across all deployments. The resulting set of light fixtures after eliminating degenerate light fixtures is output as the final set of light fixtures LF^i in the room containing light sensor i .

5.3 Experimental setup

In this section, we discuss details of our underlying sensing and ground truth hardware, and their deployment in homes, to evaluate LightSense.

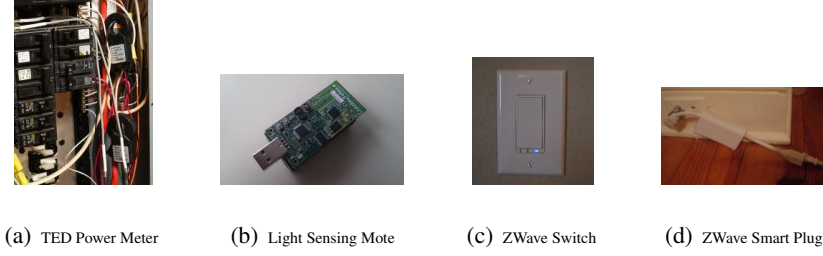


Figure 5.5: Sensors used in LightSense and its evaluation. Our system uses (a) a smart power meter on the electrical mains, and (b) light sensors in each room. For experimental ground truth, we also deployed (c) Z-wave light switches for in-line fixtures, and (d) Z-wave plug load metering for plug-in fixtures.

5.3.1 Real World Deployments

To evaluate the LightSense system, we deploy it in four multi-resident homes for 10 days each. We use the TED power meter [75] along with the cheap Hamamatsu photo diodes [76] interfaced with the telosb mote [77], as shown in figure 5.5a and figure 5.5b. The power meter samples at 1Hz to match the capabilities of prevalent smart utility meters, while the light sensors sample only at 2Hz to have a minimal impact on the lifetime of the sensing node.

A key challenge in evaluating electrical load sensing systems, as observed in the literature [30, 19], is accurately determining ground truth usage times. To measure ground truth, we replace the existing switches in each home with wireless ZWave smart switches [70], and connect each lamp to a wireless ZWave smart plug [70] that measures the power consumption of the appliance plugged into it. Figure 5.6 shows details of the four homes in which LightSense and the ground truth sensing system were deployed. We see that each home has a high proportion of rooms with windows, which introduces higher noise levels due to natural light changes for our LightSense approach to handle. During our deployments, we also encountered multiple light fixture types such as incandescent bulbs, CFLs, and halogen lights, and both dimmable and non-dimmable switches with different light intensities and wattages as seen in figure 5.8.

5.3.2 Energy Usage Ground Truth

Before we can evaluate the energy feedback accuracy of LightSense, we first need to map the set of ground truth light fixtures GF monitored using the ZWave devices, to the set of light fixtures LF inferred by LightSense. To perform this mapping, we first compute a match recall metric $mr_{g,l}$

House #	# of residents	# of rooms	# of light fixtures	House Type	# of rooms with windows
1	3	8	12	3 floor House	8
2	3	6	6	3 bedroom student apartment	6
3	2	8	9	1 bedroom condo	6
4	4	9	14	3 floor house	7

Figure 5.6: Details of the four homes chosen for deploying LightSense

between each fixture $g \in GF$ and each fixture $l \in LF$, that denotes the proportion of ON-OFF events from g that are *detected* by at least one ON-OFF event from l . An event x from fixture g is said to be *detected* by an event y from fixture l if both the ON and OFF times of events x and y are within $dT_m = 10$ seconds of each other. We map each ground truth fixture g to an unmapped LightSense fixture $map(g) \in LF$ with the *maximum match recall metric*, if the maximum match recall metric $mr_{g,max(g)} > .5$, i.e. at least 50% of the ON-OFF events from fixture g are detected from fixture l . At the end of this matching process, some light fixtures in GF and LF may be unmapped, resulting in false negatives and false positives, respectively.

To provide an easily understandable economic interpretation of light fixture energy consumption, we transform the raw energy in KWh for each light fixture to a projected energy cost over 5 years, assuming a cost of \$0.15 per KWh, and extrapolating our 10 day sample period to a 5 year period. We use the ZWave ON and OFF times to compute the energy cost of each ground truth light fixture. We set the threshold parameters dL (in raw ADC units), and dP (in watts), to both be 10 in order to identify even low intensity light fixtures; we perform a sensitivity analysis on these parameters in section 5.5.3.

5.4 Assessing Energy Usage

Given the above experimental setup, we assess the accuracy with which LightSense infers the energy usage of individual light fixtures. Figure 5.7 shows the projected energy cost of each light fixture for




Figure 5.7: Energy Costs at \$0.15 per KWh projected over 5 years, for individual light fixtures across all four homes, as computed by LightSense and the ground truth system. LightSense accurately reports the energy costs of the top energy consumers in each home.

the ground truth approach and LightSense, across all four houses over the 10 day deployment period. For each light fixture in figure 5.7, figure 5.8 shows the corresponding room locations, wattages, and light intensity changes observed at our light sensors.

We see from figure 5.7 that our LightSense approach accurately computes the energy costs of the top energy consumers in each home; we observe an average accuracy of 91.1% in determining the energy cost of light fixtures consuming 90% of the home’s lighting energy. All of the false negatives and positives have very low energy consumption and do not significantly interfere with decision making about which light fixtures to optimize. Interestingly, House 1, which is a large 3 floor house, uses less lighting energy than House 2, which is a 3 bedroom student apartment; the reason is that House 1 has large wall-sized windows, which negate the need for lighting during most of the day. Our feedback can also be used to apportion the lighting bill based on energy usage; for example, the resident in Bedroom1 of House 2 might be motivated to reduce her high lighting energy given a personal cost incentive to do so.

For each light fixture, LightSense uses Tier III to provide users with the (1) nominal wattage, (2) room location based on the light sensor room location, and (3) the total energy cost. LightSense reports the nominal wattage of each light fixture shown in figure 5.8 within $\pm 5W$ accuracy for all the top 90% energy consuming light fixtures in each home. Using these three inputs, residents

House 1				House 2				House 3				House 4			
Light #	Room	Wattage	Light change	Light #	Room	Wattage	Light change	Light #	Room	Wattage	Light change	Light #	Room	Wattage	Light change
1	MasterBed	135	23	1	Livingroom	185	366	1	Livingroom	95	155	1	Frontroom	55	311
2	Livingroom	35	193	2	Bedroom1	90	55	2	Livingroom	115	59	2	Basement	325	353
3	Livingroom	40	29	3	Kitchen	120	71	3	Bathroom2	220	1187	3	Kitchen	250	60
4	Kitchen	90	20	4	Bedroom2	35	50	4	Bedroom	95	70	4	Bathroom	395	1028
5	MasterBath	20	36	5	Bedroom3	50	414	5	Kitchen	110	181	5	Kitchen	280	125
6	MidBathroom	10	62	6	Bathroom	50	260	6	Officerroom	90	55	6	Livingroom	80	50
7	BottomBath	40	41					7	Bathroom1	305	1157	7	Dining Room	200	110
8	MasterBath	42	0					8	Diningroom	200	182	8	Bedroom	95	69
9	BottomBath	40	27					9	Livingroom	55	239	9	Bathroom1	95	649
10	Kitchen	5	-					10*	Diningroom	100	19	10	Nursery	55	129
11	TopRoom	80	923									11	Bedroom	60	32
12	Kitchen	98	-									12	Kitchen	70	-
												13	Kitchen	30	-
												14	Kitchen	110	-

Figure 5.8: Nominal Wattages and Room Locations for the light fixtures in the four houses along with the light intensity increase or decrease observed at our light sensor as the fixtures are switched ON or OFF, respectively. The light fixture numbers here match the light fixture numbers in figure 5.7.




Figure 5.9: Projected cost savings based on LightSense recommendations and an optimal system with ground truth data.

can optimize high energy light fixtures by replacing incandescent bulbs with CFLs or LED bulbs, by replacing traditional lights with motion activated lights, by implementing a daylight harvesting system, or just by using the lights in a more conservative fashion. Each of these techniques requires a different input cost and reduces energy by a different amount, so the energy feedback from LightSense is valuable to the end user in order to maximize net cost savings.

Given a limit n on the number of light fixtures that users can optimize, we compute the *maximum projected energy cost savings* achievable over 5 years (the typical lifetime of a CFL), assuming we optimize the top n light fixtures based on the energy cost ordering computed by LightSense; we assume 75% energy reduction for each light fixture after optimization, which is similar to the energy savings obtained on switching from incandescent lights to CFLs. Figure 5.9 shows the maximum projected cost savings of LightSense for each house as n is increased along the x-axis. We also show the cost savings for an optimal system that uses the ground truth energy costs. We see that LightSense closely tracks the optimal recommendation system for optimizing light fixtures across all houses. Also, in each house, the cost savings are quite different and become very small after 6-7 fixture are replaced; this suggests that users should carefully weigh the cost vs. benefit of optimizing light fixtures, before deciding on a suitable n for each house.


We see the value of LightSense by combining the results from figures 5.7, 5.8 and 5.9 ; many of the low energy, low return on investment light fixtures in each house are in similar rooms as the top energy consumers, and in some cases even have a higher or similar nominal wattage compared to the top energy consumers. Thus, optimizing light fixtures based on their actual energy usage is more profitable compared to using only their nominal wattages or room locations.

5.5 Analysis


In this section, we analyze the accuracy of LightSense in identifying individual light fixture ON and OFF events, and provide detailed insights on the high accuracy achieved by LightSense in inferring the energy cost of individual light fixtures in the home. Firstly, we show that both the data fusion and Bayesian matching steps contribute equally to aggressively filter out false positives. Secondly, we show that our Bayesian Clustering approach is needed to accurately assign power measurements to ON-OFF events. Finally, we perform a sensitivity analysis of LightSense to the power and light thresholds dP and dL .

5.5.1 Impact of Data Fusion and Matching

To understand the impact of Data Fusion and Bayesian Matching on improving LightSense accuracy, we implement five distinct inference approaches with increasing sophistication, namely: Light Edge



(a) Edge Detection Recall



(b) Edge Detection Precision

Figure 5.10: Edge Detection Recall and Precision for LightSense and its individual inference components. LightSense outperforms the individual components by effectively fusing the noise filtering properties of data fusion and ON-OFF event matching.

Detection (**LE**), Power Edge Detection (**PE**), Light Edge Matching (**LM**), Power Edge Matching (**PM**), and finally Light Power Data Fusion (**LP**), that eliminates temporally isolated light or power

edges from by applying our temporal intersection step. We also consider the **LightSense** approach, which effectively uses all the inference algorithms and sensor sources used in the five approaches above. To measure the accuracy of each inference approach above, we use the common evaluation metrics of edge detection recall and precision for light fixture ON and OFF edges. These two evaluation metrics represent the accuracy of LightSense in identifying individual light fixture events, and are different from the energy accuracy presented in section 5.4.

Figure 5.10 shows the recall and precision metrics for all approaches; we fix the light and power thresholds $dL = dP = 10$, to be low enough to detect the low intensity fixtures seen in figure 5.8. We see that approaches **LM** and **LP** improve precision significantly compared to the simpler approaches **LE** and **PE** in each house, by eliminating unmatched or temporally isolated false positive edges. Interestingly, approach **PM** only has a marginal improvement over approach **PE** due to the large number of false positive power edges. Finally, we see that in each house, *LightSense improves upon both LM and LP, by combining both noise elimination techniques to aggressively filter out false positives from the light and power meter data streams.* We also observe that in most houses, both data fusion and bayesian matching contribute equally to eliminate false positive light and power events that do not originate from light fixtures.

5.5.2 Impact of Bayesian Clustering

To understand the impact of using Bayesian Clustering in accurately assigning power costs to ON-OFF events, we consider two different power assignment approaches: (1) *Power Only assignment* uses only the power meter data to assign power edges to ON-OFF events, i.e. by only using $vprob_{x,y}^i$ in equation 2, (2) *Bayesian Clustering assignment* also uses the joint probability of observing the light and power meter data as shown in equation 2 using our lpbayes Bayesnet. For each approach, we consider the total energy cost error, normalized as a percentage by the total energy cost of all fixtures in the home. In figure 5.11, we observe that leveraging the long term temporal correlation between the light and power meter data reduces the total normalized energy cost error from 40-50% to less than 10% in multiple homes. We see almost no impact of our Bayesian Clustering approach in House 2, because this small apartment has a little noise on the power lines due to the lack of HVAC, television, or other complex appliances.




Figure 5.11: Relative Total Cost Error for power only assignment and lpbayes power assignment.

5.5.3 Parameter Sensitivity Analysis




Figure 5.12: Recall Precision tradeoff for LightSense and its individual inference components as the light and power level thresholds are decreased from 500 to 10 in steps of 50.

The main variable parameters in our approach that affect classification accuracy are the light and power edge thresholds dL (in ADC units) and dP (in Watts). Figure 5.12 shows the recall-precision graph for all the five additional inference approaches from section 5.5.1 and LightSense as dL and dP are *simultaneously* decreased from 500 to 10 in steps of 50 each; we set the minimum values

of the threshold parameters to be 10 instead of 0, and make the last step size 40. We plot the edge detection recall and precision for each parameter setting averaged across the four homes.

As the threshold parameters are decreased, we observe that (1) edge recall for all approaches increases as low intensity events greater than the decreasing threshold are detected, and (2) precision decreases as more false positives are introduced from low intensity edges. While approaches **LM** and **LP** achieve only about a 5-10% *worse F1 score* compared to LightSense at a threshold setting of 50, they achieve much lower precision when detecting low intensity events at less than 50 ADC light units or less than 50 W. However, identifying low intensity events is *important*, since we see from figure 5.8 that several light fixtures in the top 5 or 6 consumers in each home have either a wattage less than or close to 50W, or a light intensity less than 50 ADC units. Thus, our LightSense approach combines both noise elimination techniques, matching and data fusion, to accurately identify even important low intensity light fixture events.

5.6 Limitations

Although our LightSense approach achieves high accuracy in inferring the usage and energy consumption of individual light fixtures, there are several limitations in the approach and evaluation. LightSense currently does not provide feedback to the end user about the location of light fixtures detected in rooms. Therefore, if a room has multiple 100W light fixtures, extensions to the LightSense approach are needed to provide users with the locations of the light fixtures relative to the light sensor, based on their average light intensity at the sensor. Also, in larger rooms, or rooms with multiple virtual partitions, multiple light sensors are needed in the same room to ensure that all light fixtures are detected by at least one light sensor. In such a case, a methodology is needed to compute the locations of the unique light fixtures detected by each light sensor, and provide this input to the end user effectively. For dimmable lights, LightSense only reports the wattage at the typical dim level, and not the nominal wattage of the light bulb in the fixture; this makes the problem of locating reported light fixtures in the room more challenging. One solution to address the challenge of locating sensors is to provide an interactive interface to the user, that allows her to turn on and off a particular light under consideration, and enhance the LightSense approach to map the current

light under consideration to a light fixture whose usage energy consumption profile has already been created by LightSense using long term sensor data.

Another key limitation is that LightSense cannot differentiate fixtures that have the same wattage and the same light intensity at the light sensor; although such a case is rare, it is possible if a light sensor is placed equidistant to two light fixtures of the same model and wattage. Also, LightSense cannot accurately recognize the usage of dimmable switches, if they are set to a different dim level each time they are used; this is because a consistent light and power signature that enables LightSense to identify the source light fixture is no longer generated for such dynamic dimmable light fixtures. LightSense is also unlikely to achieve high accuracy with light fixtures that have a wattage of less than 10W due to its dependence on the power meter, which has very high noise in this low wattage range. Finally, LightSense is not currently set up to address multi-state light fixtures, and can only effectively infer the usage of two state light fixtures, which are the most common case; during our four home deployments, we did not encounter any multi-state non-dimmable light switches.

Improvements are also possible to our experimental setup used to evaluate LightSense in our four sample homes. The main limitation of our current evaluation is intermittent data loss from both the light sensors and the whole house power meter during the deployment period. Data loss from the light sensors is due to unreliable wireless communication, especially in larger homes; to address this challenge, more reliable packet delivery protocols need to be used. Data loss from the whole house power meter is due to lossy powerline communication used by commercial power meters [75], affected by noise on the power line and the use of electronics in the home. To address this challenge, it is necessary to use more reliable power meters that use wireline or reliable wireless communication to communicate instantaneous power usage data to the smart home system.

5.7 Conclusions

In this chapter, we presented the LightSense system to infer the fine-grained energy usage of light fixtures in homes, without requiring direct sensors on individual light fixtures in the home. Through 10-day deployments in four natural home environments, we show that LightSense effectively combines two noisy sensor sources, namely light sensors with 25% precision and whole house power consumption data with 15% precision, to infer the energy costs of individual light fixtures with

91.1% accuracy. Our LightSense and WaterSense approaches address the challenging problem of recognizing the use of individual water fixtures and electrical fixtures that are ambiguous in their smart meter profiles. Both approaches use a simple, cheap sensor in each room, and do not require any configuration or training effort from the end user.

Our LightSense inference algorithm requires an edge detection tier to detect unique events from the power meter and light sensor data, similar to the WaterSense algorithm. However, unlike the AUTOLABEL and WaterSense algorithms, room clustering of light sensors with motion sensors is not a primary concern, since we assume that the motion sensors and light sensors are present in the same wireless sensor device. Rather, a key challenge is to eliminate false positives in both the light sensor data and power meter data; this challenge is absent in both the AUTOLABEL and WaterSense approaches because of the inherent nature of the binary sensor and water flow data. To address this new challenge, we developed effective data fusion techniques in Tier II of the LightSense algorithm to combine the noisy light sensor and power meter data streams to eliminate noisy events.

Our data fusion and Bayesian clustering techniques from LightSense, can be used to combine *any* pair of sensor streams whose noise is not temporally correlated, in order to detect distinct physical phenomena that affect both sensor streams in a consistent manner. Similar to WaterSense, we use a Bayesian matching approach to match positive events with negative events in the sensor data, corresponding fixture ON and fixture OFF events; however, unlike WaterSense, we find that the Bayesian matching in LightSense turns out to be an effective tool to eliminate a number of false positives from both the light sensor and power meter data streams.

By combining the WaterSense and LightSense approaches with our AUTOLABEL inference approach, we expect that future activity recognition systems will be able to effectively reduce both sensor installation effort, and any configuration or training effort, in using home activity recognition systems; in other words, our vision is that users should be able to purchase a handful of sensors, deploy them in the home with minimal effort, and start using the turn-key activity recognition system for medical and sustainability applications.

Chapter 6

Non-Invasive Resident Identification

In earlier chapters, we presented sensing and inference solutions to reduce user effort in configuring and installing accurate activity recognition systems in the home. However, these activity recognition solutions are only able to infer which activities are being performed, and not who performed them. These single person activity recognition solutions are useful in many scenarios. For example, in single person homes with an independent elderly resident, such solutions could be used by commercial peace of mind monitoring companies such as QuietCare and Wellaware [7, 8], to provide relatives of the elderly and caregivers with the elderly resident’s daily activity record. Such solutions may also be useful in whole house intelligent heating or cooling solutions such as the smart thermostat approach [5], which uses only the *aggregate* activities of all residents in homes.

However, in order to respond to activities in the home, many applications require that the activity recognition system must be able to identify *who* is in a particular location or performing a particular action such as cooking or exercising; thus, the ability to *identify residents* is crucial for many activity recognition applications in multi-person homes. For example, in sustainability applications, it is important to know who is using more energy or water so that utility bills may be apportioned correctly, or who is in a particular room so that the temperature or light levels can be adjusted to improve comfort for the residents in the home. As another example, in medical applications, caregivers are interested in *who* is unable to sleep well, or is visiting the toilet an abnormally large number of times.

Existing innovative implementations that perform resident identification in homes have several advantages, but also have drawbacks. Some approaches are inconvenient because they require the user to wear a tag [31, 32], or to manually trigger a biometric sensor such as a thumbprint or retina

scanner [78]. Some systems require cameras for gait, form, or face recognition [35], but cameras are often perceived as invasive because they can be used to collect much more information than just the user's identity [36]. Other implementations require structural changes to the home, such as instrumenting the floor [79, 80] with force plates, which can incur high cost and effort. Many practical smart home applications such as in-home medical care for the elderly [81, 8] and occupant-based energy monitoring [15] cannot use solutions that inconvenience the user, are intrusive, or require an expensive building retrofit.

Our recent discussions with the WellAware elderly monitoring company [7] reveal several interesting user requirements for accurate, long term elderly resident identification and tracking in homes: (1) residents will not wear tags or manually identify themselves at every room for long periods of time, (2) residents will not allow perceived invasive devices such as cameras or microphones in the home, and (3) residents want the sensors to be fairly invisible, similar to existing motion sensor installations, and do not want an expensive building retrofit. Since existing implementations have some drawbacks with respect to the above requirements, commercial deployments by the elderly monitoring enterprise [7] today are limited to single-resident homes or do not fully monitor information about multiple residents.

In this chapter, we examine the use of biometric height sensors to satisfy the above requirements for both the elderly monitoring enterprise [7], and a wide variety of other activity recognition applications. Height sensors have several advantages over existing approaches: they are cheap, convenient and minimally invasive for the residents, and not very time consuming to install in an existing home. We use ultrasonic distance sensors mounted above the doorways to measure the height of individuals that walk through the doorway. The inherent accuracy of height sensing is too low for reliable biometric identification from a large population of individuals: in controlled experiments in the lab, height sensors requires a **7cm** difference in height to differentiate people with 99% accuracy, and most people have heights within a small range from 160-180 cm. However, we make two key insights that allow height to be an effective biometric sensor *in the home*. First, most homes have very few residents: height may be a weak biometric for differentiating between 20 or more people, but is likely to be very effective in homes that have only 2-4 residents. Second, people move through a home in predictable ways, as determined by the floor layout: if height sensors are

placed above every doorway, then the history of height measurements can be used to potentially surpass the inherent accuracy of the sensor.

In this study, we evaluate the use of height as a biometric in four ways: (1) We quantify the biometric error of our approach using 20 subjects in a controlled laboratory environment, in which we vary the direction, speed, and location of the person walking under the doorway. (2) We perform a first pass evaluation of using height events to differentiate residents at doorways as they enter and leave rooms in 3 homes for 5 days each (3) We use public anthropometric data containing the heights of elderly residents in 2077 multi-resident homes from the 2006 health and retirement study to estimate that our approach can potentially achieve at least 95% identification accuracy in **85%** of elderly US homes sampled in this study, based on extrapolating the accuracy results from our controlled experiments in the lab. (4) Through a simulation study, we show that incorporating the history of height measurements at multiple doorways using a tracking approach can potentially increase the proportion of homes where our solution is applicable with 95% accuracy from 85% to 95%, and also reduce the height difference required for 99% identification accuracy from 7cm to 3.25cm. We quantitatively compare our approach against two other state of the art non-invasive resident identification implementations, namely anonymous binary sensor and activity model based multi-resident tracking [50], and weight sensing [80], and find that our approach achieves improvements in identification accuracy or installation effort and cost compared to existing implementations of these approaches.

In this chapter, we present the first ever analysis of biometric height sensing for home activity recognition applications. A key limitation of our study is that we do not characterize the proportion of false positive and false negative height sensor events in real home environments, and also do not develop a holistic approach to track residents' locations in homes based on the height sensor data streams. As mentioned in chapter 1, based on our key insights on height sensors, there has been significant follow up research on designing and evaluating complete tracking solutions for in-situ home environments using height sensors. In contrast to that work, we lay the groundwork in this study that shows the potential of using weak biometric sensors in residential monitoring scenarios; we do not address the challenge of continuous tracking of residents in in-situ scenarios. The main challenge as observed in the follow up research is to address false positives and negatives caused in the height sensor data due to complex home environments and multiple height sensors being placed

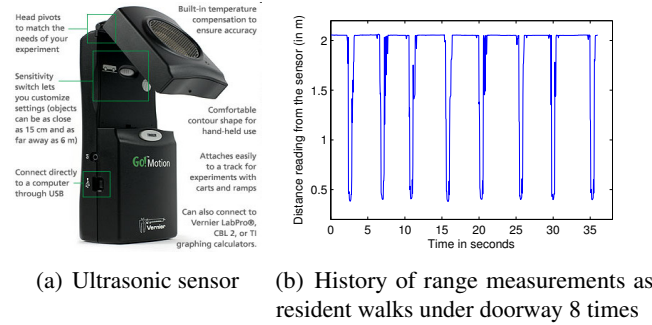


Figure 6.1: Sensor and sample data used in our resident identification study. Our study used the Go Motion ultrasonic range finder mounted above doorways (a). As users walked beneath the sensor, the range measurements changed (b).

very close to each other. In addition to removing false positives and negatives, it is also necessary to determine the direction of residents as they walk through the doorway in order to accurately track residents in the home. Finally, another limitation of our ultrasonic height sensing approach is that it causes annoyance to pets in the home, and thus cannot be used in homes with pets.

The rest of this chapter is organized as follows. Section 6.1 discusses existing resident identification solutions from the literature in more detail. Section 6.2 gives an overview of our approach and describes our algorithm to sense height from ultrasonic distance measurements. Sections 6.3 and 6.4 describe the results of our controlled lab experiments and preliminary in-home experiments respectively. Section 6.5 analyzes how our empirical results potentially extrapolate to a national level using public anthropometric data. Section 6.6 presents the simulation study to evaluate the potential improvement in the inherent identification accuracy of the height sensor by tracking the history of height measurements at multiple locations in a home.

6.1 State of the art

Resident identification in smart homes is a long-standing problem with many existing solutions. In this section, we discuss a representative sample of these solutions, including their advantages and disadvantages, and their applicability with respect to our user requirements.

Tag-and-track approaches operate by placing a uniquely identifiable device on each individual resident. This approach has been widely used since the Active Badge system almost two decades ago [33], and in other systems since. The pedestrian localization system proposed by Woodman et al

[31] uses a foot mounted inertial sensor to track pedestrians to within .7 meters 95% of the time in a large office building with no additional infrastructure. More recently, innovative tracking solutions that require a very low infrastructure cost [82, 83] are emerging. Tag and track approaches have three important advantages: (a) High location granularity with little or no infrastructure, (b) Selective preservation of location privacy by switching off device, and (c) Highly scalable with respect to the number of residents in an indoor space: each user can be given a new device with a uniquely identifying number. However, one drawback of this approach is that it requires the user to actively carry the device at all times when location information is desired. It can be an inconvenience to in-home residents for long-term deployments. In our past experience with deployments, and while trying to use MoteTrack [32] wireless tags to collect ground truth in this study, users frequently forget to carry their tag, especially immediately after waking up or showering. Automatically reminding residents to carry the device is an option, but we believe such an approach is intrusive and inconveniences the user for long-term applications such as elderly medical monitoring.

Several indoor resident identification systems use **cameras** for computer-vision based face, shape, and gait recognition [35], or **microphones** and audio signal processing. These approaches might require expensive on-board computation or high communication bandwidth to a central base station that executes the vision algorithms, but are passive and highly accurate. However, user studies by researchers from Intel and companies like WellAware have found that a large fraction of potential users have perceived privacy concerns about cameras or microphone sensors [36]. Therefore, this class of approaches is most appropriate for short-term situations in which rapid deployment, and/or high accuracy are important, and where long-term privacy concerns of monitoring residents in their own homes are not an issue.

Wilson et al, in 2005, propose using only resident **usage models of anonymous motion sensors in rooms and switch sensors on daily-use objects**, and **resident activity models** to identify and track their activities and locations [50]. They propose using a particle filter that uses Markov state transition and sensor use models learned from short term training data, obtained using a tag and track approach or manual labeling. The main advantage of this approach is that the simple single-pixel sensors are cheap and easy to install, and are not perceived to be invasive or inconvenient. However, an important drawback of this approach is low accuracy: this system was reported to have 70% accuracy when tracking 3 residents over a week-long period, and in our own deployments in 3

multi-resident homes, we observed this approach to have accuracies of 65-75%. These accuracy rates may be reasonable for some applications, but confusing the identities of residents more than a third of the time could cause problems for some smart home applications such as medical monitoring. To increase identification accuracy, additional biometric sensor data, as discussed in this chapter, can be included in the STAR particle filter.

Several systems including the Active Floor and the Smart Floor **instrument the floor** to locate and identify individuals [80, 79]. Jenkins et al, in 2007, studied the effectiveness of using resident mass, derived using force plate signals, to identify multiple individuals in a large population [80]. Gait analysis can also be used to differentiate individuals from instrumented floors. This type of single-pixel mass-based identification approach has the advantage that it can be performed without inconveniencing the user or violating resident privacy. Existing force plates and smart floors require careful installation to improve aesthetic appeal and achieve user acceptance; accurate force plates are also very expensive. However, more compact, cheap designs of weight sensors that have the same form factor as a floor mat can be explored for easier installation and better aesthetic appeal in the home.

Height is a weak biometric that is often used on driver's licenses or police reports, and it can not be used to definitively identify individuals from a large population . Some existing systems use invasive video cameras to identify height [84]. Nishida et al [85] propose instrumenting the entire ceiling with a dense set of ultrasonic devices to perform fine-grained location tracking with an ultrasonic radar system. However, this system is not evaluated for its ability to differentiate or identify individuals, and this approach would involve substantial deployment effort. In 2006, Jenkins et al [86] proposes placing infrared or ultrasonic distance detectors on top of doorways for identification based on height (in a poster), using an approach similar to that described in this chapter. However, height sensors are not experimentally evaluated for accuracy, multiple readings are not combined as the user walks through the home to improve accuracy, and the poster does not analyze the wider ramifications of height sensing on *in-home* resident identification. To the best of our knowledge, our study is the first to analyze how height sensing can potentially be used to effectively address the multi-resident room location and identification problem in homes with high accuracy.

6.2 Overview: Sensing Height with Ultrasonic Sensors

To identify residents as they move throughout a home, we deploy an ultrasonic distance sensor above every exit and entry into a room. We used an off-the-shelf ultrasonic distance sensor [87] shown in figure 6.1a. This distance sensor sends out ultrasonic pulses at 50KHz in a diverging cone 15 to 20° off the axis of the beam. The device then measures the time taken for the echo to return, and uses it to calculate the distance to any obstacle in front of it. Only the minimal distance of any obstacle is reported. For example, when a resident stands under the device, only the distance to the top of the head is reported, while distances to the shoulder, ear etc are automatically filtered out.

Figure 6.1b shows the example data from the distance sensor as a subject walks repeatedly under the ultrasonic sensor mounted on top of a typical doorway eight times. The **default distance** reported is 2.1m, which is the distance when there is no obstacle in front of the ultrasonic sensor. When a subject walks under the sensor, we see minimal peaks that correspond to the ultrasonic beam making contact with the subject's head; the difference between this minimal distance and the default distance of 2.1m returns the *apparent* height of the person as she walks under the doorway. In our controlled experiments, described in the next section, we observe that this *apparent* height measured while walking, is on average less than the *erect* height measured while standing, by about 1-3cm.

Our algorithm to extract height events and height values is as follows. We first compute timestamps when the reported distance is below the default distance with no obstacles. We then cluster these timestamps using the DB-SCAN clustering algorithm [61] to compute discrete **height events**, that correspond to residents passing by or standing under the sensor. This clustering process eliminates most noise due to a single, spurious reading. Then, for each cluster of low readings, we find the minimum distance reported (i.e. the maximum height value measured). We subtract that measurement from the default height measurement with no obstacles and use the result to be the height measurement for that height event.

To *identify* residents based on measured height values, we use a Maximum Likelihood Estimate (MLE) classifier to assign each height event to one of multiple candidate residents in a home. For each height event, the MLE classifier simply computes the probability that each resident triggered it, based on the height of that resident and the error distribution of the sensor, and assigns the height event to the resident that maximizes the likelihood of the observed measurements. In the next section,

we collect height data from 20 test subjects using controlled experiments in a lab to characterize the error distribution of height measurements, under diverse scenarios of passing through a doorway.

6.3 Experiments in a Controlled Lab Environment

6.3.1 Experimental setup

We characterize the error distribution of height sensors in a controlled lab setting by placing the ultrasonic sensor on top of a doorway about 90 cm wide and having 20 users with known heights pass beneath the sensor in a controlled manner. We chose a doorway of this width because it matched the width of many of the doorways seen in our real home deployments. We selected 20 subjects of differing heights for our experiment. The distribution of heights among the 20 subjects can be inferred from any of the scatter plots in figure 6.2. The subjects were randomly chosen from a pool of graduate students from 20-30 years of age; 16 of our subjects were male and 4 were female.

For each subject, we first manually measured the height while standing using a tape measure. We then measured the height reported by the ultrasonic sensor when the subject stands still exactly under the sensor. The subject then walked under the sensor several times as we varied the configurations of our doorway and requested changes in the direction and speed of walking. In particular, every subject (1) Walked 20 times in a simulated narrow doorway measuring 75 cm in width, (2) Walked 21 times under the full doorway 90 cm in width (7 times perpendicular to the plane of the doorway, and 7 times on two perpendicular planes at an angle of 45° to the plane of the doorway, for a total of 21 times). We repeated the above experiments with and without shoes for each resident.

6.3.2 Evaluation results

Figure 6.2a illustrates that, when residents are standing erect beneath the sensor, the average error across all 20 subjects is only 0.2cm, and the maximum height error is 1.15 cm. Figures 6.2b and c show scatter plot of mean height measurement error and standard deviation in error while *walking*. *Error* here refers to the difference between the manually measured height and the height output by our height based identification algorithm for each height event. The results shown in figure 6.2 use the aggregated height data from all our walking experiments without shoes. We do not include results

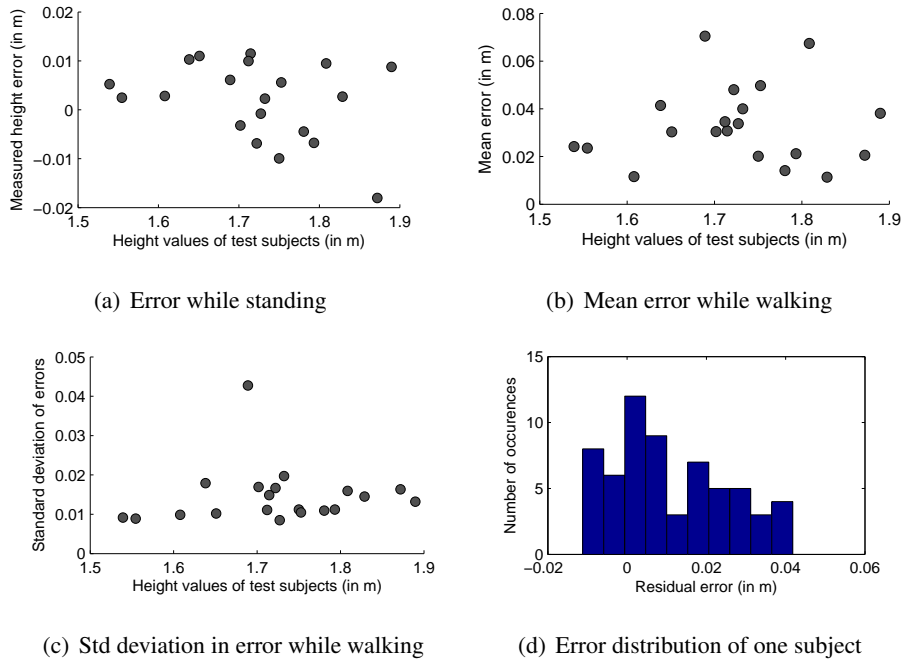


Figure 6.2: Experimental results for height based identification in the lab. Controlled laboratory experiments indicate low measured error while standing. Mean error while walking is higher due to a natural reduction in height compared to standing erect, and different walking styles. Standard error while walking is very low. The error distribution approximates a log normal distribution.

with shoes here, but observe that the mean measured height simply increased by the height of the shoe on average, and changes to the standard deviation of errors were negligible with shoes on.

From figure 6.2b, walking height as measured by our sensors is lower than erect, standing height by 3.31 cm on average across all subjects. This is possibly due to the natural decrease in *apparent height* as a person walks. Also, different walking styles such as bending and keeping heads down contribute to this decrease compared to erect standing height. More important is the *deviation* in residual error, the standard error, for each subject across different height events, since this will be crucial in determining identification accuracy among multiple residents in a home. We note that the mean deviation in error is only 1.45 cm. This low deviation implies that 99% identification accuracy can be obtained as long as the heights of two residents are 7cm apart. We explore this tradeoff more fully in section 6.6, when we describe our history based tracking algorithm using height.

In figure 6.2d, we show an example distribution of residual error from one test subject. The distribution shown here suggests a log normal distribution rather than a normal distribution for residual error. Thus, we ran hypotheses tests at .03 significance level for all subjects to test two




Figure 6.3: Identification Accuracy as the number of individuals to be differentiated is increased. Height measurements become less effective for biometric identification as the pool of individuals increases. The heights in our study were easier to differentiate than those of the general population.

different hypotheses (a) distribution of height values is normal (b) distribution of height values is log-normal. *The proportion of subjects for which the normal and log-normal hypotheses could not be rejected are 75% and 85% respectively.* The log-normal distribution, which skews naturally to the right, appears to be a better fit for modeling measured height. This is because the apparent height of a person very rarely *increases* (perhaps due to thick shoes) while walking, but more often decreases in experiments due to the ultrasonic beam making contact with the person's side (e.g.) shoulder or ear) instead of the head.

Using the empirical height data collected in the lab, we compute how well height can differentiate among a fixed set of N residents in an indoor space. In particular, we empirically calculate the accuracy with which height events are assigned to their ground truth test subjects. In figure 6.3, we show how this empirical identification accuracy using height decreases as we increase the number of residents under consideration in the indoor space; we randomly choose residents from our pool of 20 subjects, and evaluate how accurately individual height events generated by our subjects are labeled using a log normal MLE classifier trained from the controlled experiment data. For each N value, we repeat the random sampling 100 times. Also shown in figure 6.3 is how this analysis extrapolates to a national level. We model the mean and standard deviation in residual height error as a function of height using two types of curve fit models: simple linear curve fitting, and nearest neighbor interpolation.

Thus, given the height of a resident, we can derive his/her mean and standard deviation, and use this in turn to derive the mean and deviation of the corresponding log normal distribution of measured height.

The 2006 health and retirement study [88] (HRS 2006) contains height measures of 4154 elderly residents living in multi-resident households. We randomly sample a fixed number N from this set of residents to be identified in an indoor space. We then *analytically* calculate the probability that any height event is assigned to the correct resident among the N residents, assuming a MLE classifier that uses the log normal distributions derived from our nearest neighbor and linear fit models. In figure 6.3, we show how this probability of correct event labeling degrades as we increase the number of residents in the home, randomly sampling 100 different sets of elderly residents for each N value. We only show the results with the nearest neighbor curve fit model, since only negligible differences were observed when the linear fit model was used.



Figure 6.4: Height sensors deployed above doorways in a home.

As we can see in figure 6.3, for indoor spaces strongly resembling 2 or 3-resident elderly homes, height based identification has a mean accuracy of 87-92% using both empirical data and extrapolation to the national level. In particular, we note that 99% of the elderly multi-resident households with valid height measurements in the national study were 2-resident households. We use this insight in the next section to demonstrate the high accuracy of height based identification in 3 real multi-resident home deployments. For households with 4 residents, the identification accuracy drops to 77%; in section 6.6, we show that by using the history of height events at multiple doorways in a home, we can improve the identification accuracy in even 4-resident homes to 90%.

6.4 Experiments in Home Environments

Our controlled experiments characterize the sensitivity of height measurements to various conditions, but these experiments do not reveal the frequency with which these conditions occur in a home environment. To evaluate the accuracy of height based identification in such an environment, we deployed ultrasonic sensors in three homes for five days each. The ultrasonic sensors were deployed




Figure 6.5: Sensors used in our height based identification study in homes. Motetrack tags and beacons (left) were used to collect ground truth locations. Motion and magnetic reed switch sensors (middle/right) were used to evaluate STAR.

on doorways of rooms similar to our controlled experiments, and can be seen in figure 6.4. In addition to the ultrasonic sensors, we deployed the motetrack indoor localization system [32] in all homes to get ground truth locations of residents. Motetrack is a tag and track approach to localization that requires each resident to carry a mote. It uses trained RSSI signatures from beacon nodes, like the one shown in figure 6.5a, to localize the mobile motes in the home.

The main goal of the natural home deployments is to evaluate the ability of height sensors to identify residents at doorways in the home, *and* compare it to a state of the art non-invasive multi-resident identification solution that *only* relies on simple activity models derived from labeled binary sensor data [50]. In order to make this comparison, we also deployed anonymous X10 motion sensors in every room, and X10 switch sensors on daily-use objects such as the fridge, microwave, stove etc. Figures 6.5b and c show examples of the motion and switch sensors used in the homes. In section 6.5, we also compare with another well studied non-invasive resident identification solution [80] that uses resident mass to differentiate between residents in a home.

Table 6.1 shows some of the deployment details for the 3 homes, including number of rooms, and ground truth height values of the couple living in each home. Given the large differences in height values in the three homes, we expect our height based identification solution to perform with high accuracy. We evaluate the accuracy with which biometric height sensors are able to identify residents in a home. We compare the accuracy of our approach with a state of the art passive identification technique based only on 'biometrics' of simple activity models of residents, derived from labeled binary sensor data, as evaluated by Wilson et al in their STAR approach [50].

First, we temporally cluster X10 motion sensor firings from the same room using db-scan [61], to identify discrete room visits of residents in the home; we assume here that these temporal clusters correspond reliably to ground truth room visits of residents. Ground truth resident labels for the

temporal clusters are obtained from motetrack’s location trace. Our aim is to assign resident labels to each of these clustered room visits, using either biometric height sensors, or using the location trace for each resident computed in STAR using only activity models of residents. To assign resident labels to room visits using only height sensors, we run the log normal MLE classifier on each height event that occurred during the temporal cluster. When the MLE classifier assigns a height event to a resident, that resident is added to the list of labels for that temporal cluster.

The STAR resident tracking system proposed by Wilson et al [50], uses individual Markov state transition and sensor observation models of residents to track their activities and locations. The essence of their tracking approach is that individual residents have different movement/activity patterns in the home, and/or have unique sensor use patterns. Similiar to their original implementation [50], we simply restrict our state space to include current room location of individual residents. The state transition and sensor observation probabilities are learned using counting from our ground truth training data obtained from the motetrack location trace; we performed leave-one-out cross validation over the 5 days of room occupancy data obtained from each home, i.e. for each day, we tracked residents using Markov models trained from all the other days’ data. We implemented a multi-hypotheses tracking solution to track room visits of multiple residents in the home, similar to the particle filter solution implemented by Wilson et al [50].

Home	Number of rooms	Height of resident A in m	Height of resident B in m
1	7	1.88	1.77
2	4	1.68	1.55
3	5	1.75	1.63

Table 6.1: Details of the 3 homes used in deployments to evaluate height based resident identification

Figure 6.4 compares the identification accuracy of our height sensor approach and the existing approach based on activity and binary sensor use models. We see that identification based on simple activity models and binary sensor use models only achieves accuracy around 65-75% in 3 homes, while height based identification achieves accuracies ranging from 98-100%. Clearly, the activity and *binary* sensor use patterns of residents in these homes are not distinguishing enough to assign room visits with high accuracy. We do not claim here that our approach is better than STAR; instead, we simply compare with an existing instantiation of the STAR framework using only ‘biometric’ room

transition models and binary sensor use models. Certainly, as pointed out by Wilson et al [50], by using more fine-grained sensing at a higher installation cost than our approach, it is possible to better differentiate residents even using these simple models; one could even incorporate height sensor data in the STAR particle filter.

We also observe that our height sensor based approach does not require any training phase, while any approach that depends mainly on activity or binary sensor use models requires a long training phase where ground truth locations of residents need to be collected using wearable tags, to determine the probability models used in tracking; such a training phase might also require the installation of a separate infrastructure just for tracking, although recent advances in low cost tag and track solutions [83, 31] may negate the need for such a tracking infrastructure.




Figure 6.6: Accuracy of height based identification and activity model based identification. Height sensors achieve higher accuracy than achieved by existing implementations of the STAR approach that use only activity models and binary sensor use models of residents for identification.

6.5 Accuracy of Height Sensing in Homes Nationwide

We analyze the proportion of homes at a national level where our height solution can *potentially* differentiate residents with high accuracy using the 2006 health and retirement study (HRS 2006) [88], which contains height and weight measures of elderly residents living together in the same household. Of the 2107 multi-resident households with valid height and weight measures for every resident in the home, we used the 2077 households that were two-person households. We do not have currently have access to any anthropometric datasets that support our claim for a wider population demographic. We also note here that wider, longer term deployments in real homes are the best way

to evaluate this technology, and our results below are best effort extrapolations from our controlled experiments in the lab.

For each home, using the height values of the residents in the home, we first derive a log normal probability model for each resident in the home using his/her height and the curve fit models described in section 6.3.2. We then *analytically* calculate the probability that any height event will be assigned to the right resident, assuming that each resident is equally likely to generate a height event. From now on, we refer to this probability as the **probability of correct resident identification** in a home. For each home, HRS 2006 also provides the weight measures of every resident. Jenkins et al [80] in 2007 observe that weight based identification using force plates has a Gaussian error with a mean of 0.67kg and standard deviation of 0.96. Assuming this Gaussian model and mean parameters, we calculate the probability that any gait event will be assigned to the right resident, assuming again that each resident is equally likely to generate a gait event.

Given the probability of correct resident identification for each home in the sample, we compute the proportion of homes where the probability of correct identification is above a fixed threshold; Figure 6.7 shows how this proportion decreases as we increase the threshold for probability of correct resident identification. Our height based identification solution is potentially applicable to 85% elderly homes in the US with at least 95% identification accuracy. Using force plates and weight based identification, up to 92% of the elderly homes can potentially achieve least 95% identification accuracy. Given the high cost and effort involved in retrofitting a home with force plates, height based identification is preferable, even though it is slightly less accurate; however, alternative cheaper sensing solutions for weight measurement or gait analysis can be explored for preferential use over height sensors in some homes.

6.6 Improving Height Measurement Accuracy with History

We have shown in the previous sections that height based biometric resident identification is potentially applicable to a significant proportion of elderly homes in the US with high accuracy. However, from the analysis seen in figure 6.7, height sensors achieve *less* than 95% identification accuracy in 15% of the homes. In this section, we show how information such as the room topology of a home,




Figure 6.7: Applicability of height based identification at a nationwide scale. Height sensors are potentially applicable with high accuracy to a large proportion of US elderly households. Weight sensors have potentially higher applicability, but require good sensor design to aesthetically install on the floor.

and the past history of height sensor events on multiple doorways in a home, can potentially be used to improve the inherent accuracy of biometric identification using the height sensor.

As an illustrative example, assume two residents A and B initially in the bedroom. Assume that after some time, resident A leaves the bedroom, and goes to the kitchen through the living room to get a snack. Even if a few *individual* height sensor events lead to incorrect results from the MLE classifier, the *sequence* of height events generated by A will have a higher likelihood of being classified as resident A; we use spatio-temporal continuity of motion through the constrained floor layout of a home to improve identification accuracy. An assumption in the in the example above, and in the analysis below, is that the error at individual height sensors is independent; this assumption may be true most of the time. However, there might occur cases where the error is more systematic, such as a person stooping over to carry a heavy object; in such cases, the utility of using the sequence of height measurements could potentially be reduced.

We use a simulation based study that is driven by the public-use height data from HRS 2006 and our height error models derived in section 6.3.2, to estimate the improvement in identification accuracy that can be achieved using the history of height events in a home. We assume a 6 room home across all the elderly households for consistency. We define the same HMM for each resident's room transitions, to indicate equal transition likelihood from one room to another; we do this to ensure that differing room transition patterns of residents or specific room topologies of homes do




Figure 6.8: Potential improvement in height based identification accuracy from using a tracking approach. Our simulation study shows that the history of height measurements collected over the track of a resident through the home potentially improves identification accuracy and applicability of height sensing in US elderly homes.

not unfairly improve the identification accuracy possible by using the past history of height events alone. We generate 1000 height events for each home in HRS 2006, using our HMM to generate room transitions, and using the height error models from section 6.3.2 to generate noisy height events for each resident. We assume a height sensor at every entry/exit into a room.

Given the simulation trace, we evaluate two approaches to identify resident labels for height events in the home (1) A naive MLE classifier that only considers data from individual height events (2) A probabilistic multi-hypotheses tracker that uses past history and room topology, embedded in a HMM Viterbi approach. Figure 6.8 shows the applicability of height based identification across elderly homes in the US with and without history information, based on the results of our simulation experiment. When our probabilistic multi-hypotheses tracker is used, we observe that height based identification can potentially achieve at least 95% identification accuracy in 95% of elderly homes in US, as opposed to only 85% of elderly homes covered by the naive identification algorithm.

Figure 6.9(a) provides more insight into the scenarios where history information might be most useful. When the height difference of the residents living in the home is small, using the past history of height events greatly improves the accuracy over naive MLE classification. Using tracking history can potentially reduce the height difference required for 99% identification accuracy from 6.9cm to 3.25 cm. Figure 6.9(b) demonstrates another important benefit of using history, the ability to achieve higher accuracy in indoor spaces with more residents; the heights of residents in the hypothetical

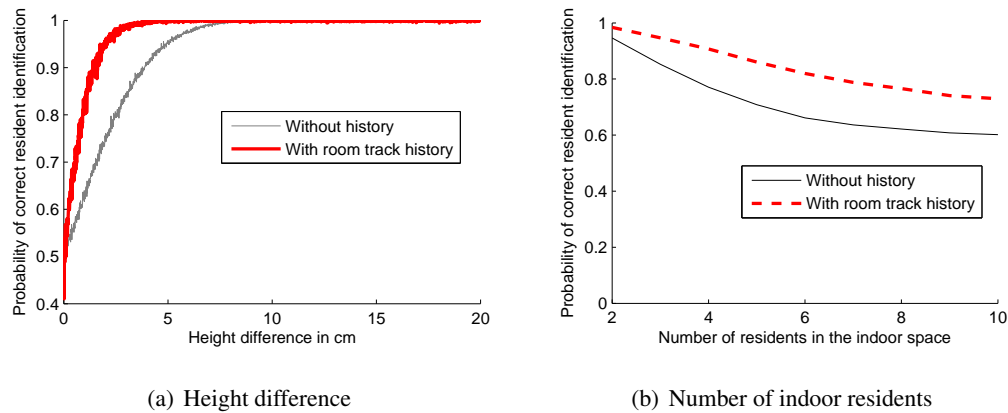


Figure 6.9: Impact of tracking on height difference required for accurate identification, and the number of individuals that can be differentiated accurately. A history of readings could potentially decrease the height difference required for accurate identification(a), and increase the number of residents that can be reliably differentiated(b)

multi-resident homes (100 sample homes at each point) are randomly generated using height data from the 4154 elderly residents from HRS 2006. By using the history of height events in a home, we can potentially improve the identification accuracy in 4 person homes from 77% to 90%.

6.7 Limitations

Although our novel height sensing solution shows promise in non-invasive resident identification for home activity recognition applications, there are several limitations with our approach and evaluation that could be addressed in future research. A key limitation of our approach is that our in-situ evaluation of height sensing accuracy does not consider false positives and negatives in individual height sensor events due to the lack of accurate ground truth; rather, motion sensor events are used as an approximation of room visits to be labeled with concurrent height sensor events. Better ground truth solutions need to be developed to label each height event; this is a challenging problem, since cameras might be unacceptable for long term resident labeling, while wearable radio tags pose accuracy problems in homes for ground truth purposes. False positives and negatives, and height sensor inaccuracy, are possible due to multi-path effects from ultrasonic sensors, especially in (1) rooms with wide doorways that require several adjacent height sensors to achieve sufficient coverage, and (2) adjacent doorways very close to each other. The multi-path challenge could be addressed using a combination of a distributed synchronization algorithm, and careful placement of the sensors.

Another promising approach is to add simple infrared range sensors, that suffer less from multi-path effects but are also less accurate, to eliminate uncorrelated noise in the ultrasonic height sensor data.

Another key limitation is that we do not develop and evaluate a complete tracking approach to estimate the residents' current room locations based on the height sensor data. Our simulation study to evaluate the improvement in height sensor accuracy from tracking makes two simplified assumptions that are challenging to handle in a holistic tracking approach: (1) we do not add false positives and negatives to height sensor events, which occur in practice, (2) we assume that we know the direction of room transition for each height sensor event. Recent research has shown that handling these two challenges introduces additional state in the tracking algorithm, and a HMM based Viterbi algorithm is no longer able to handle the large solution space; to handle the computational complexity, particle filter algorithms are being developed to track the likely resident states and height sensor data associations. Another limitation of our study is that we have not considered the challenges in height sensor based tracking posed by guests with similar heights, multiple shoes, pets, or residents with gait based disabilities.

An inherent limitation of our approach is that it only identifies residents at doorways, and not inside rooms; such an approach is suitable for room level localization of residents. In contrast, camera based techniques [35], tag-and-track approaches [31, 32], and recent solutions that use a dense array ultrasonic rangefinders on the ceilings [85], are able to track residents at a meter level granularity. The lack of meter level location of residents could potentially limit the accuracy in assigning individual, concurrent activities of two residents in the same room. To address this challenge, additional height sensors could be deployed, or variations of the data association algorithm for simultaneous tracking and activity recognition, proposed by Wilson et al [50] could be explored. Another limitation of using ultrasonic sensors is that it disturbs pets such as cats and dogs in the home.

Another limitation of our study is that we extrapolate our empirical results in the lab environment to a nationwide data set which is restricted only to *elderly* multi-resident home; we restricted ourselves to the elderly demographic, since anthropometric measures for other populations were not publicly available. An interesting extension would be to explore how our solution generalizes to a larger population, including young couples, small and large families, and multi-resident student homes, by conducting large scale surveys of anthropometric measures in these homes. Since our approach is based on a resident biometric, it cannot be applied in all homes with high accuracy,

unlike existing approaches such as tag and track. For applications that require higher identification accuracy than offered by height sensing alone in a given home, we propose to explore adding multiple non-invasive sensing modalities including floor mat sensor implementations for weight measurement, color sensors above the doorway, and potentially additional infrared or ultrasonic range sensors to measure the profile of a person as they walk through the doorway.

6.8 Conclusions

In this study, we analyze the use of ultrasonic range sensors placed above doorways in a home to identify residents for multi-person activity recognition systems. Height is typically a weak biometric, but we make two key insights that make it promising for *in-home* monitoring. First, height is highly effective among small populations where the height differences among residents are likely to be large enough for reliable differentiation. Second, residents walk through the home in predictable, constrained patterns dictated by the floor layout, and the multiple height measurements of the resident as they walk through multiple doorways in the home can be potentially be used to improve the inherent accuracy of the height sensor. In this study, we quantify the error with which ultrasonic height sensors measure the heights of residents as they walk under the doorway. Using publicly available height measures of residents from multi-resident elderly homes and the height error distributions derived from our controlled lab experiments, we extrapolate that a resident identification accuracy of at least 95% can potentially be achieved in 85% of elderly homes using a naive classification algorithm and in 95% of elderly homes using an approach that aggregates measurements from multiple doorways.

Height measurements can potentially be used for accurate resident room level tracking and activity recognition, if height sensors are placed above every entrance and exit to a room, as proposed in this paper. A key challenge in achieving high room level localization accuracy is to address the problem of false positives and negatives in the ultrasonic height sensor data. Of course, any biometric sensors including thumbprint or retina scanners [78], could be placed at the entrance of any room to locate residents, but violates our requirement of not requiring manual identification effort from residents. Tag and Track systems such as Pedestrian localization [31] and Motetrack [32] can be inconvenient to the user, STAR [50] *using activity models* has low accuracy and requires an inconvenient training phase, and weight sensing using force plates or smart floors [80] to track

residents would require a costly installation. Of these existing implementations, our height sensor approach shows significant promise as a cheap, convenient, and non-invasive solution for identifying residents in homes.

Chapter 7

Wireless privacy for home activity recognition systems

In the previous chapters, we presented several solutions for designing convenient activity recognition systems for the home, without requiring high sensing or configuration cost and effort. An important consideration overlooked so far in this dissertation is user privacy in home activity recognition systems. As mentioned in chapter 1, Activities of Daily Living (ADLs) are typically very personal and private, and must be kept secret from third parties. In particular, activity recognition systems deployed for remote medical monitoring applications, are obliged by HIPAA regulations [40] to protect this ADL information from unauthorized users.

In all of the activity recognition systems that we have presented so far, *wireless* data collection is preferred to using wireline networks to collect data from sensors distributed in the home. Wireless sensors are much easier to install, especially in a home environment, where it is intrusive and time consuming to place wires connecting each sensor to a central base station. As seen in our recent guide on home sensing deployments, alternative power line based communication is also not feasible in the home environment [30]. Indeed, wireless home sensor systems have already been deployed at a large scale in US homes. Over 32 million homes in the US have wireless security sensors installed on doors and windows, and motion sensors installed inside and/or outside the home [41]. Over 5 million homes have wireless X10 devices [65] and ZigBee devices such as wireless doorbells, appliance controls, wireless smoke detectors, and wireless light switches. These sensors produce a constant

record of the *Activities of Daily Living* (ADLs) within a home, thereby enabling applications such as elderly monitoring [28], energy conservation, and home security [65].

Although wireless communication greatly eases installation effort and cost for home sensing systems, it does pose a privacy risk, since the data is broadcast in a public medium, that is potentially accessible to an adversary outside the home overhearing the wireless transmissions with a snooping device. To protect the ADL data from an adversary snooping outside the home, wireless sensor data is usually encrypted using one of several encryption systems developed for wireless sensor networks [89, 90].

In this chapter, we present a powerful new attack that allows us to observe private activities in the home such as cooking, showering, toileting, and sleeping by snooping on the wireless transmissions of sensors in a home, *even if all of the transmissions are fully encrypted*. This attack needs only the timestamp and the *fingerprint* of each radio message, where a *wireless fingerprint* is a set of features of a RF waveform that are unique to a particular transmitter. Thus, we call this the *Fingerprint And Timing-based Snooping* (FATS) attack. Our FATS attack uses the unsupervised AUTOLABEL algorithm discussed in chapter 3, to transform the timestamps and fingerprints of wireless transmissions, to ADL data for exploitation by an adversary.

Our FATS attack is a new attack that has not previously been studied or demonstrated. Wireless fingerprinting and traffic timing analysis are both well-studied techniques. However, we are the first to combine these two techniques to create and demonstrate a novel privacy attack in the kind of single-hop, wireless networks common in residential ubiquitous computing. In this chapter, *we present solutions for designing wireless activity recognition systems in residential environments to preserve private activity information about the residents from FATS attacks*. For example, we hide packet transmissions from the adversary using signal attenuators, we introduce random delays on transmissions, and we generate spurious transmissions in order to decrease the effectiveness of the FATS attack. We find that each of these techniques has a different tradeoff in terms of privacy protection performance, hardware costs, and application/user costs. Based on the cost-benefit profile of each technique, we present guidelines for applying them most effectively in typical home activity recognition systems. We evaluate these design guidelines and show how many of these guidelines can be used together in a hybrid fashion to yield very high privacy protection with minimal implementation costs.

We empirically evaluate both the FATS attack and our suite of privacy preservation techniques by using the sensor data from the activity recognition systems in eight different homes used to evaluate our AUTOLABEL approach. Our results suggest that the FATS attack is highly effective even without knowing any prior information about the home, achieving around 80-95% accuracy on activity recognition in many homes in the best case. The design guidelines that we propose greatly decrease the effectiveness of the FATS attack, reducing the inference accuracy to anywhere from 0 to 15%, while greatly reducing implementation costs.

The rest of this chapter is organized as follows. In section 7.1, we present related wireless privacy attacks, and show how they are different from our proposed FATS attack. In section 7.2, we briefly discuss how an adversary uses our AUTOLABEL inference algorithm to infer private activities of residents in the home, and evaluate how such an attack performs in a realistic environment, containing both packet loss and errors in wireless fingerprinting. In section 7.3, we present a suite of privacy solutions to the FATS attack, and show how many of these solutions can be used together in a hybrid fashion to protect against the FATS attack in a home environment.

7.1 Background

FATS is a type of *side channel attack*, which means that it uses information revealed by a cryptographic system other than the ciphertext to infer either the cryptographic keys or the original data [91]. Some well-known side channel attacks include the TEMPEST attack, which uses leaked electromagnetic radiation from a computer monitor to infer the plain text input to the system [92], and a recent study that exploits physical properties of variable bitrate encoding schemes to infer the movie a person is watching on commercial devices [93].

Wireless fingerprinting is a well-studied technique in which physical characteristics of RF transmissions are used to differentiate between messages from different radios, even when those radios have the same model and manufacturer. Statistical features of transient signal *amplitude* have been used to fingerprint Bluetooth devices [94] with false positive rates of 5% and detection accuracies as high as 93%. Similar results have been observed on 802.11 WiFi radios [43] and on *sensor nodes* using the ChipCon CC1000 radio [44]. We discuss the hardware requirements for fingerprinting when we discuss deployment details in the next section. The above work on *wireless*

fingerprinting is different from *software-based fingerprinting*, such as recent work that identifies and tracks users of 802.11 enabled devices by exploiting implicit identifiers in 802.11 network traffic [95]. Wireless fingerprinting is usually used to *enhance* privacy by enabling hardware-based authentication [43], not to compromise it. A recent study proposed using wireless fingerprints to compromise privacy in vehicular sensor networks [96], but did not combine fingerprints with traffic timing analysis.

The FATS attack is different from most existing traffic analysis attacks. Previous work has demonstrated that Internet traffic patterns in wired networks can be used to match a sender with a recipient [97], and multi-hop radio traffic in wireless sensor networks can be used to locate the sensor source or the base station [98, 99]. The countermeasures for these attacks require changing network flow patterns at the routing level. The FATS attack uses fingerprints to do traffic analysis in a single-hop network, and so is not affected by such countermeasures. Yang et al [100] describe a related traffic analysis attack in which an adversary can infer when a network event has occurred by observing global transmission timestamps alone, but this work does not combine transmission timestamps with wireless fingerprints.

Our FATS attack infers activities such as cooking, showering and toileting. It simply leverages the AUTOLABEL algorithm discussed from chapter 3 to infer private ADLs of residents in the home. The FATS attack assumes a home fitted with simple wireless sensors [101] from the invisible man approach to activity recognition, as discussed in chapter 2. Our FATS attack would be less successful with other sensing approaches to activity recognition, that use RFID tags on household objects coupled with wearable RFID readers [54]. Our attack is more suited to a *wireless sensor system* than to an *RFID-based system*, since it's easier to snoop on relatively long range wireless transmissions than on short-range RFID signals. As we have argued in this thesis, the wireless sensor approach to activity recognition is more convenient to the end user, and thus, we believe that our attack is an important consideration to take into account when building future residential activity recognition systems.

7.2 The FATS Attack

The FATS attack adversary uses a wireless snooping device to record the timestamps and unique source fingerprints of each wireless transmission. This series of timestamps and source identifiers, is input to our AUTOLABEL activity inference algorithm, discussed in detail in chapter 3. We do not repeat the algorithm and its evaluation here, but note that our three tier AUTOLABEL algorithm is first able to infer coarse grained activities such as sleeping and home occupancy from Tier 0 using just the transmission timestamps, as discussed in section 3.1 in chapter 3. The FATS attack is then able to combine the source fingerprints with the timestamps in Tiers I, II and III, to automatically infer the room visits of residents to rooms such as the bathroom and kitchen, and also private ADLs such as cooking, showering, toileting, and eating hot/cold food. Since our AUTOLABEL algorithm does not require any prior information about the home or its residents, it can be easily used by our FATS adversary to infer private ADLs of residents.

To fingerprint radio sources, the FATS adversary could either simply use the RSSI values [102] of radio messages, which is less accurate, or use an antenna connected to a high frequency oscilloscope or a commercial integrated signal analyzer with built-in fingerprinting software [103]. Even though many of these snooping devices are currently costly and power hungry, it is expected that cheaper devices will be available as wireless fingerprinting matures and becomes increasingly important for hardware-based authentication. It is easier for an adversary to launch the FATS attack if she has access to uninterrupted power, either from an unattended outdoor power outlet nearby, an adjacent apartment or a distant surveillance area with powerful directional antennae. We have performed experiments to validate that wireless mica motes inside our office buildings can be snooped upon from the outside. Also, X10 radios have a long radio range and snooping on these devices is as easy as driving around with an X10 receiver to receive unencrypted X10 camera data [104].

As mentioned earlier in this chapter, we evaluate our FATS attack using the sensor data from eight homes, as presented in chapter 3 to evaluate the AUTOLABEL activity recognition approach. It must be noted that all of this wireless sensor data is recorded by a base station inside the home, that understands the source identity of each message. We use this data to *emulate* an adversary outside the home snooping on the timestamps and fingerprints of wireless transmissions. In particular, we emulate the adversary first assuming a best case scenario, where all packets are received, and

fingerprinted correctly to indicate the source of the wireless transmissions. We then emulate the adversary using a realistic setting, introducing both packet loss and fingerprinting errors in the sensor data collected at the base station in each home.

To measure how accurately an adversary can infer the private ADLs of residents in the home, we use the same evaluation metrics used in chapter 3, and discussed in section 3.2.2, namely Event Detection Rate (EDR), True Positive Rate (TPR), and Duration Accuracy (DA). In the best case scenario, we observe that the adversary receives the same data received at the base station in the home, and thus achieves the same recognition accuracy as our AUTOLABEL approach from chapter 3, as discussed in section 3.3. Thus, the FATS adversary is able to infer Tier 0 activities such as sleeping and home occupancy with 85-100% accuracy, Tier II room visits with 85-90% accuracy, and Tier III activities such as cooking, toileting, and showering, with around 80% accuracy, on par with supervised activity recognition approaches in the literature.

We also evaluate the FATS attack in practical, non-ideal settings, where the adversary will observe both packet loss and fingerprint errors. The evaluation in figure 7.1 models packet loss as a percentage of the nodes hidden from the adversary. Even under pessimistic scenarios such as 30% of the nodes hidden from the adversary, the FATS attack still infers about 50% of Tier II and Tier III events and detects Tier 0 events with about 80% DA. Tier 0 DA is unaffected by fingerprint errors, since it simply relies on the transmission timestamps. Figure 7.5 shows how fingerprint errors, modelled as a feature deviation in the x-axis and explained in detail in the next section, affect Tier II and Tier III EDR. Under pessimistic feature deviations of about 50% of the total feature space (deviations of 20 for a feature space of 40), the adversary still infers upto 50% of Tier II and Tier III events in the home. Assuming both 20% packet loss and 20% feature deviations, the adversary still infers Tier II and Tier III events with 40% EDR, and Tier 0 events with 85% DA on average. Thus, even under pessimistic scenarios, the FATS attack is effective in recognizing the private ADLs of residents in the home, and this motivates the need for effective privacy solutions to reduce the inference accuracy of the adversary to less than 10 or 15%.

7.3 Design Guidelines to Enhance Privacy

In this section, we evaluate a number of techniques to thwart the FATS attack in future home activity recognition systems. Each of these techniques has a different cost-benefit profile, and we analyze how and when each technique is most effective. We present the following five guidelines for building wireless ubiquitous computing systems in homes or residential environments. We justify each guideline in subsequent subsections. We conclude by presenting a **hybrid approach** that combines many of these guidelines to achieve very high privacy protection with very low implementation costs.

1. **Signal attenuators** should be deployed in a select few rooms such as kitchens or bathrooms where many activities occur to effectively mask activities in these rooms.
2. **Random delays** of the order of 15 to 20 minutes should be added to the transmissions of sensors in the bathroom and kitchen that are involved in short duration activities to effectively hide these activities.
3. **Periodic transmissions** should be used on binary or low bandwidth sensors that are typically involved in long-duration activities, such as bedroom and living room sensors.
4. **Fingerprint masking** should be used on time-critical sensors like fall detection sensors, where latency introduced by random delays or periodic transmissions is unacceptable, or on sensors/rooms where signal attenuators are infeasible.
5. **Spurious or fake transmissions** should be combined with real transmissions for sensors such as camera or microphone sensors that cannot afford the high energy cost from periodic transmissions.

7.3.1 Using Signal Attenuators

The most obvious approach to protect against the FATS attack is to prevent the adversary from hearing messages in the first place. There are several possible signal attenuators, and we list three here: (i) Using very low power transmissions and a multi-hop route to the base station: this incurs a moderate hardware cost in terms of additional router nodes, and reduces reliability (ii) Using a wired connection to the base station: this requires quite a lot of deployment effort and time, and (iii) Using Faraday cages: this is expensive to set up and prevents outside communication from hidden rooms, which is essential for eldercare applications. We recommend scheme (i) since it has the least cost among the various schemes, though it does not have the same protection guarantees as scheme (ii).




Figure 7.1: Effect of Signal Attenuators on event detection rate (EDR) of Tier II and III events, and on duration accuracy (DA) of Tier 0 events

We implement two schemes for signal attenuators in our evaluation, namely: (i) **SAA** - Signal Attenuators in All rooms, and (ii) **SABK** - Signal Attenuators in the *Bathroom* and *Kitchen* only, as per design guideline #1. Figure 7.1 shows the effectiveness of the FATS attack as we increase the percentage of nodes hidden by signal attenuators from 0-90%. Figure 7.1(a) shows that signal attenuators are very effective at reducing the EDR for Tier~II and Tier~III activities; with 40% hidden nodes, EDR is reduced to about 30-35% in scheme SAA and to 20% in scheme SABK. Thus, many of the activities detected in Tiers II and III can be hidden effectively by using signal attenuators in just the bathroom and kitchen. Figure 7.3 shows that signal attenuators do not have a strong effect in reducing the Tier~II and Tier~III DA and TPR. The TPR actually increases as more nodes are hidden, because the number of events detected becomes small enough that almost no events are spurious. Figure 7.1(b) shows that signal attenuators, even if implemented in the entire home (SAA), are not very effective at reducing the DA for Tier 0 information; with 40% hidden nodes, the DA is nearly 90% in single person homes, and is over 60% in multi-person homes. This is because the sensors that are not hidden by the signal attenuators still reveal the *presence* of human activity very well. Thus, we recommend in design guideline #1 that signal attenuators be used in select rooms such as the bathroom and kitchen (SABK) where many activities occur to reduce implementation cost and hide these activities effectively.




Figure 7.2: Effect of random delays on event detection rate (EDR) of Tier II and III events, and on duration accuracy (DA) of Tier 0 events

7.3.2 Using Random Delays

Because all tiers of the FATS algorithm rely on transmission timestamps, modifying the transmission time is one way of reducing the effectiveness of the attack. Thus, we propose adding random delays to sensor transmissions bounded by a maximum delay parameter. We implemented two schemes in our evaluation of random delays: (i) **RDA** - *Random Delays on All* sensors, and (ii) **RDBK** - *Random Delays on Bathroom and Kitchen* sensors involved in short duration activities only, such as cooking and toileting. Figure 7.2 shows that the effectiveness of the FATS attack decreases as increasingly long random delays are added to sensor transmissions. Figure 7.2(a) shows that even small random delays of about 10 minutes are highly effective at reducing the EDR for Tier III events such as cooking and showering to about 50-60% under both schemes RDA and RDBK. This is because random delays introduce errors in Tiers I, II and III which use the device transmission timestamps as input. Also, figure 7.3 shows that random delays of about 30 minutes reduces both DA and TPR of Tier II and III events to about 40%; i.e., more than half the events detected by the adversary are false positives, and the inferred duration of these events is highly inaccurate. As seen in figure 7.2(b), longer random delays of about 3 hours reduce Tier 0 DA to about 60% in single person homes and a much lower 25% in multi-person homes. In multi-person homes, there are only a few hours per day with no activity from *all* residents, and 3 hour random delays ensure that almost no period is inactive, resulting in the larger duration errors. However, we do not recommend such long delays of the order of hours to hide long duration Tier 0 activities such as home presence and sleeping.

Despite their effectiveness, random delays conflict with the requirements of real-time sensors such as wireless light switches, or fall detection sensors that need to transmit data immediately. Also, we need to consider if random delays of the order of 15-30 minutes are acceptable to the end users; they might be certainly acceptable to remote healthcare providers who are only interested in long term trends such as a decline in the ability of residents to perform ADLs. For other users, we need user queries to verify if such delays are acceptable. Given the similar performance of schemes RDA and RDBK in protecting Tier III events, we recommend in design guideline #2 that random delays of the order of 15 minutes be applied to non-emergency sensors in the bathroom and kitchen if the end users find such delays acceptable.

7.3.3 Using Periodic Transmission

If a sensor transmits periodically instead of only when it has data to transmit, it makes the transmission timestamps independent of the data, making it impossible for the adversary to infer any information. Thus, periodic transmission on all sensors guarantees 100% privacy. We first estimate the extra power consumed by periodic transmissions for typical *binary home sensors*, using empirical power consumption data on the telos mote with the CC2420 radio [105], a popular hardware platform in wireless sensor network research. We assume that a mote sleeps for a latency period P and wakes up to transmit a large enough data payload to capture *binary event information* for the past P seconds, along with the mandatory header fields. The total percent reduction in node lifetime for different periods P is shown in Figure 7.4. As the period P increases, the percent reduction in node lifetime decreases along with the power consumed by periodic transmissions; this is because as the period increases, the number of packets and the associated energy wastage from packet overhead decreases. We note that *the total reduction in node lifetime for periodic transmission with a period of 8 seconds is only 8.75% of the total original lifetime of the node*; thus, periodic transmissions is an excellent solution for binary sensors. The low power consumption here is because the network is one-hop, and nodes do not need to go into receive mode, unlike other sensor network deployments.

Despite their effectiveness, periodic transmissions cannot be applied to (i) real-time sensors, because of the delay limitation, and (ii) high bandwidth sensors, because of the excessive power consumed by redundant transmissions. Thus, we recommend in design guideline #3 that periodic transmissions be used for binary, non-emergency sensors in the living room and bedroom. This will




Figure 7.3: Effect of signal attenuators, random delays, and fingerprint masking on duration accuracy (DA) and true positive rate (TPR) of Tier II and III events




Figure 7.4: Effect of period of transmission on node lifetime

hide Tier 0 DA with only a negligible power cost, where the other techniques are expensive. Indeed, we implemented periodic transmissions on bedroom and living room sensors and found that this reduced Tier 0 DA by 17% in single person homes and a significant 71% in multi-person homes.

7.3.4 Using Fingerprint Masking

We can also preserve privacy through *fingerprint masking*, in which we hide the fingerprint of a transmitter. One approach to doing this might be to use potentiometers instead of resistors for all radio circuitry, and to vary these during each transmission. Another approach might be to wire together multiple sensors and use a common radio for all of them, so that the individual source fingerprint is hidden. As a variation, a sensor can also be wired to multiple radios, each of which might have a different set of sensors assigned to it, to further obfuscate the fingerprints. One problem with fingerprint masking is that it creates an *arms race* scenario in which the adversary and hardware designer must continually try to outsmart each other to uncover and hide new features respectively. Because of this challenge and the need to change existing radio hardware, we recommend in design guideline #4 to use fingerprint masking only in sensors that cannot tolerate any delays, such as wireless light switches or fall sensors and in sensors/rooms where signal attenuators are infeasible.




Figure 7.5: Effect of fingerprint masking on event detection rate (EDR) of Tier II and III events

For evaluation purposes, we use a simple model to simulate various degrees of fingerprinting error. We assign a scalar fingerprint ID_i using a uniform random distribution to each device i such that $0 < ID_i < L$. When a device transmits, the adversary observes some fingerprint $\hat{ID} = \mathcal{N}(ID_i, \sigma)$ and identifies the transmitter to be $\text{argmin}_j |\hat{ID} - ID_j|$. Thus, fingerprinting errors are likely to increase as the standard deviation σ increases. We simulate fingerprint errors on the raw data by gradually increasing σ until it equals the actual length of the feature space L , set to 40 in our case. Figure 7.5 shows that small fingerprint errors are effective at reducing Tier II and III EDR. Also, figure 7.3 shows that fingerprint masking has a similar effect to random delays in terms of Tier II and III DA and TPR, causing a significant drop in both duration accuracy and true positive rate. Introducing fingerprint errors is effective because it distorts sensor clustering as devices from different rooms appear to fire together, and also distorts the features used in our classifiers.

7.3.5 Introducing Spurious or Fake Transmissions

Yang et al [100] propose a countermeasure for traffic timing analysis in which fake and real transmissions are combined in such a way that a fixed probability distribution is maintained for time between transmissions. Real transmissions are delayed to follow the probability distribution when necessary. This countermeasure can be applied to our FATS attack too. Similar to periodic transmissions, fake transmissions would essentially ensure 100% privacy with some transmission delay but with a lower power consumption, since we are only adding some fake packets rather than transmitting large constant data payloads periodically. Thus, we recommend in design guideline #5 to use spurious or

fake transmissions on high bandwidth sensors such as cameras or microphones that transmit data occasionally.




Figure 7.6: Effect of hybrid schemes on Tier II and Tier III EDR and Tier 0 DA across all homes

7.3.6 Hybrid Schemes

Based on our design guidelines, it is clear that each of our privacy solutions is best suited to certain kinds of sensors and is better at protecting certain types of information. Therefore, many of these solutions will be used in *combination* in a real home wireless system with diverse sensor types. To demonstrate that such a hybrid approach offers excellent privacy protection, we implemented several hybrid solutions that include the techniques listed in our design guidelines, namely PT - Periodic Transmissions on living and bedroom sensors, Signal Attenuators and Random Delays (SA and RD) on bathroom and kitchen sensors, and FM - Fingerprint Masking. The effect of these hybrid schemes on Tier II and Tier III EDR and Tier 0 DA is shown in figure 7.6. The hybrid schemes shown are annotated with details relating to the extent to which each technique was applied: for example, *40% FM* refers to 40% feature deviations, *RD 20 min* to a 20 minute random delay, and *40% SA* to 40% hidden nodes. Based on the results shown in figure 7.6, scheme *PT + RD 20 min* should be used effectively in homes where such delays are acceptable to hide information from all tiers of the FATS attack with minimal costs. For homes where such delays are not acceptable, one of the other hybrid schemes should be chosen based on the implementation costs affordable. Scheme *PT + 20% SA + 20% FM* looks promising, since it requires neither extensive signal attenuation nor extensive fingerprint masking, but achieves excellent privacy preservation with relatively small costs.

7.4 Limitations

While we show that our wireless fingerprint based snoop attack is an important privacy problem for future smart homes to address, there are a few limitations with our current evaluation of the attack and the associated privacy solutions. Firstly, we do not evaluate the ground truth packet reception rate, and fingerprinting accuracy, for a typical adversarial snooping device deployed outside the home; it would be interesting to measure the typical packet reception rate and fingerprinting accuracy, in order to better evaluate the exact activity recognition accuracy that could be achieved by the FATS attack. However, a key challenge in such an evaluation is that different wireless standards such as bluetooth, Wifi, and zigbee protocols may have different fingerprinting accuracies and packet reception rates for receivers outside the home. To address this challenge in our study, we emulate a range of packet reception and fingerprinting accuracies to evaluate the spectrum of activity recognition accuracies achievable by the FATS attack.

Another limitation is that we emulate many of our privacy solutions such as radio signal attenuation and fingerprint masking in our evaluation. It would be important in future research to physically implement these solutions to better understand the cost vs privacy tradeoffs involved in these solutions. Finally, we do not provide formal proofs to evaluate our privacy solutions, but rather show their effectiveness against our existing FATS attack implementation. It would be interesting as part of future research to explore whether an adversary can design more robust inference algorithms that can better adapt to packet loss, fingerprinting errors, fake messages, or random timestamp delays; it would also be interesting to develop formal models for the inference attack, and explore if we can better estimate the privacy achieved through our suite of privacy solutions.

7.5 Conclusions

In this chapter, we showed how our AUTOLABEL inference algorithm could be used by an adversary to recognize private activities of residents inside the home fitted with wireless sensors, in spite of encrypted wireless transmissions. Firstly, we show that simple activities such as home occupancy and sleeping could be inferred just from the timestamps of wireless transmissions from the home, without any knowledge of the data content or source of each wireless message. Secondly, we

show that with additional information about the source of each wireless message, through wireless fingerprinting techniques [94, 43, 44], our AUTOLABEL algorithm from chapter 1 can be used by the adversary to recognize detailed activities and room visits of residents in the home, even if all wireless transmissions are encrypted. Our FATS attack is evaluated on existing home activity recognition systems with binary sensors used in the invisible man approach, and achieves an activity recognition accuracy of 80-90% in the best case scenario for the adversary.

With emerging infrastructure mediated sensing techniques, many of the direct binary sensors in the home could be replaced by a single sensor on the home's electrical or water infrastructure that simply outputs data periodically, such as the smart meters discussed in chapters 4 and 5. Thus, the FATS attack could be deprived of detailed fixture usage information, potentially reducing the accuracy of an adversary in inferring fine-grained activities such as toileting vs washing. However, since wireless occupancy sensors are likely to be present in each room in future smart homes, the adversary would still have access to information about the individual room visits of residents, such as bathroom or kitchen visits. The adversary could potentially achieve moderate accuracy in inferring detailed activities of residents inside these rooms, by just using temporal features of these room visits, such as the start time, and duration.

Our design guidelines to protect against the FATS attack may become increasingly important as wireless sensors become more ubiquitous in homes and residential environments. Millions of homes are already vulnerable to the FATS attack, and new systems are being deployed at an ever increasing rate. Also, we believe that the FATS attack is just one instance of many potential physical-layer privacy attacks on wireless ubiquitous systems. Other attacks could be carried out in offices, factories, and even in urban-scale wireless networks. For example, a company that shares an office building with a competitor may infer a new product launch by the competitor by observing increased traffic in certain areas. This study demonstrates the power and ease of physical-layer wireless privacy attacks such as FATS, and our design guidelines are a first step toward thwarting such attacks.

Chapter 8

Conclusions

In this dissertation, we first identified three key drawbacks of existing activity recognition systems that inconvenience the end user, namely:

1. Existing activity recognition systems require users to provide extensive configuration input and training data of the order of several days.
2. Existing systems require high cost sensors, or a large number of difficult to install sensors, to provide the sensory input necessary for activity recognition algorithms.
3. Existing activity recognition systems for multi-person homes use invasive sensors such as video cameras, or inconvenient wearable sensors that need to be carried all the time.

Our goal in this dissertation was to overcome the above limitations, and design convenient, non-invasive, accurate activity recognition solutions for in-home residents. Our main hypothesis was that by effectively using data fusion techniques, leveraging the existing smart meter infrastructure in homes, and using only weak biometric sensing, we can build convenient and accurate home activity recognition solutions for the end user. We note that we are indeed able to make progress towards our goal, by leveraging our main hypothesis. In section 8.1, we discuss our main contributions towards achieving this goal, and present the advantages of our approach over existing activity recognition approaches. In section 8.2, we discuss example medical and sustainability applications that are enabled by our activity recognition system, some of which have already been implemented and demonstrated in real homes. In section 8.3, we discuss the contributions of this dissertation beyond

smart home systems. Finally, in section 8.4, we list the current limitations of our activity recognition approach, and future improvements and directions to integrate our solutions as a holistic activity recognition system for smart home environments.

8.1 Key contributions towards home activity recognition systems

In this dissertation, we have made several key contributions to designing effective home activity recognition systems. In this section, we summarize these contributions, and place them in the context of existing and potential future activity recognition approaches.

8.1.1 Lower training and configuration effort

Firstly, our novel AUTOLABEL inference approach [37] is the first unsupervised activity recognition algorithm that requires users to neither label the sensor semantics [54] nor label daily activities [28, 26]. We observe that our unsupervised activity recognition approach achieves an accuracy of 80-90%, which is similar to accuracies achieved by state of the art supervised approaches [26, 54, 28]. In contrast to existing approaches require users to label all their daily activities for 5-10 days [26, 28], our approach requires no training data about daily activities from the end user. Our approach also achieves higher accuracy than other unsupervised approaches based on mining the web for textual activity models [54] (40-60% accuracy), that require users to individually label the 40-60 sensors in the home. Our tiered inference approach clusters unknown sensors to rooms, and jointly leverages prior knowledge about room usage patterns, and sensor usage patterns, to label room locations and sensor semantics of each individual sensor in the home; from the room locations and sensor semantics of each sensor, a cross-home activity classifier is used to accurately recognize daily activities. Such an approach is preferable to existing approaches that require users to configure the activity recognition system by labeling the individual sensors and/or labeling daily activities, and is likely to be more acceptable to end users who wish to buy and use a commercial product for home activity recognition.

8.1.2 Fewer, easier to install, low cost sensors

Secondly, our LightSense and WaterSense approaches [38] eliminate the majority of sensors on water fixtures and electrical fixtures in the home, thus allowing users to more easily install the home activity

recognition system. We achieve only a 10% loss in accuracy as a result of eliminating individual direct sensors on fixtures in the home. Existing approaches to eliminate sensors on individual fixtures require configuration effort per fixture, along with a specialized expensive single sensor on the home's water or electrical infrastructure [55, 56, 57, 58]. Our hybrid approaches use an existing smart meter on the home's infrastructure along with few simple sensors distributed in the home, and are completely unsupervised; our approaches reduce both configuration and sensor installation effort per fixture. While our approaches do typically require a sensor per room, we believe that this sensor can be piggybacked on room level occupancy and identity sensors that will likely be used in future smart homes built for medical monitoring and energy conservation applications.

The Bayesian clustering technique developed as part of our WaterSense approach will be critical in future smart home systems that need to differentiate similar fixtures in multiple rooms of the home, with similar signatures derived from an infrastructure sensor [38, 57]. Our Bayesian data fusion approach developed in LightSense will be useful in effectively combining any pair of noisy sensor sources in the home with temporally uncorrelated noise, to detect phenomena that affect both sensor streams in a consistent manner.

8.1.3 More convenient resident identification

Thirdly, our height sensor based resident identification approach [39] is a significant contribution over existing inconvenient approaches that require users to carry a tag or allow camera use at all times [31, 32, 35]. Our solution is particularly useful in non-critical applications such as peace of mind elderly monitoring or home energy conservation applications. Our approach shows significant promise in early studies, and shows the potential to achieve 95% accuracy in 95% elderly homes in the US, if challenges arising from multi-path effects in homes can be overcome with sensor fusion approaches. The occupancy sensors and light sensors used in our WaterSense and LightSense approach can be integrated with our height sensors into a single doorway sensor for future smart home applications.

8.1.4 Improved wireless privacy

Finally, we discover a novel wireless snooping attack on smart homes [37], that will become increasingly important to address as wireless smart homes are commercialized and deployed in a larger scale. Our wireless snoop attack shows the potential to achieve 80-90% activity recognition accuracy in the best case, even for activities such as cooking, showering, and toileting. Our hybrid solution approach of introducing small random delays, and enforcing periodic transmissions when feasible, greatly reduces the adversarial inference accuracy of such an attack to less than 10% without a significant impact on the cost or functionality of the smart home system.

8.2 Example applications enabled by our activity recognition approach

Our proposed activity recognition system enables a number of practical applications for smart homes. In this section, we provide a few concrete examples of these applications.

Firstly, in the medical monitoring domain, monitoring daily activities of residents is useful in early diagnosis and continuous care for several ailments, such as depression, and Alzheimer's disease. For example, the Empath approach [106] geared towards depression monitoring requires information about daily activities of residents in the home such as showering, sleeping, toileting, and eating. Detection of changes in the duration, ordering, or presence of daily activities can be very useful in the continuous treatment of medical conditions such as Alzheimer's disease. An ongoing area of research is to develop effective inference algorithms and anomaly detection approaches to apply on the daily activity information inferred by our non-invasive activity recognition approach.

Secondly, in the energy conservation domain, our activity recognition approach has already been shown to provide accurate feedback to the end users about the energy consumption of individual electrical and water fixtures in the home [38]. Additionally, our non-invasive sensing approach has been used in the smart thermostat approach to leverage the sleeping, and home occupancy activities of residents to automatically control the thermostat and save energy in the home [5]. Our height sensor based identification approach could also potentially be used with emerging smart zoning approaches, to individually control the temperature in each room, or switch off heating in unoccupied rooms. In general, information about the location and activities of all residents in the home could

enable smart control of everyday appliances such as the television, lights, computer, etc., resulting in high energy savings with little effort from the end user.

8.3 Contributions beyond smart homes

Although we focus on our contributions towards smart home activity recognition systems in this thesis, many of our approaches could be extended to benefit other cyber physical systems beyond smart homes. Firstly, our unsupervised activity recognition framework could potentially be applied to wearable sensor systems such as smartphones. Two key challenges for wearable activity recognition systems is that the end user has to train the system specific to the location of the wearable sensor on the body, and variations in the location of wearable sensors require re-training. To address these challenges, a tiered inference approach could be used to first infer the body location of each wearable sensor, similar to AUTOLABEL, and then cross-person classifiers could be explored for unsupervised activity recognition across individuals.

Secondly, our data fusion approaches developed in LightSense and WaterSense could be applied to other problems in cyberphysical systems that require sensor fusion. Both data fusion approaches address the key challenge of efficiently combining instantaneous sensor data from multiple sensors with historical sensor evidence to effectively detect and classify physical events affecting multiple sensors. The data fusion approach developed in WaterSense could be applied to other sensor fusion problems that combine a distributed proximity sensor network with another single point sensor to detect events that generate both a proximity event, and a location specific noisy signal in the single point sensor; for example, one data fusion application could combine sparsely distributed sound sensors with densely distributed proximity sensors to infer the typical location of certain animals in a large forest. Similarly, our LightSense data fusion approach could be used to combine other pairs of noisy sensors with high false positive rates to accurately detect distinct physical events that generate consistent signatures in both the sensor sources.

Finally, our wireless privacy attack based on overhearing wireless fingerprints of devices could have broad implications beyond smart home systems. Other indoor wireless systems used in offices and commercial buildings could be subjected to a FATS style attack to infer protected information. For example, rival companies could learn of potentially high activity prior to an unannounced product

launch. The specific inference algorithms, and external knowledge, used in these unexplored wireless fingerprinting attacks could be different from the three tier algorithm used in our approach. Wireless fingerprinting attacks could also be applied at an urban scale to localize an individual by observing the unique wireless fingerprint of his personal smart phone or tablet device. As wireless technology becomes more prevalent, effective fingerprint masking techniques may need to be included in future wireless devices to protect against fingerprinting attacks similar to FATS.

8.4 Future improvements

Though our individual solutions to address the limitations of existing activity recognition systems show promise when evaluated in home environments, a key direction for future research is to integrate the various solutions presented. For example, the three tier AUTOLABEL algorithm is currently designed to work on a sensor system consisting of simple binary sensors. With the addition of single point sensing solutions, the use of individual fixtures such as the sink or the stove are recognized using flow or power profiles at the whole house smart meters; therefore, information about the sensors in tier III could be inferred from smart meter profiles and this sensor information could be used to constrain the room labels in tier II. Also, in our AUTOLABEL approach which uses direct sensors, a room clustering approach that uses a single distance measure between every pair of direct sensors achieves high accuracy; however, in WaterSense, to cluster noisy flow events to the motion sensor in each room, a probabilistic Bayesnet approach is more appropriate. In future systems that may contain both direct sensors and single point sensors, a hybrid clustering approach may need to be designed to effectively cluster together both direct sensors with flow or power events from the smart meters.

Another key direction for future research is building on our existing height sensor based identification solution, to evaluate it for tracking purposes in in-situ environments. Such a system requires height sensors that can sense both identity and direction, and eliminate noisy edges that may arise due to multi-path effects. Recently, a particle filter based tracking algorithm has been developed to track the locations of residents in the home, in spite of false positive events from the height sensors. A potential approach to eliminate false positives would be to combine infrared range sensors to eliminate noisy events from ultrasonic range sensors, and use the ultrasonic distance sensors to estimate identity and direction. Another key challenge in evaluating such a tracking system over a

long term is obtaining accurate ground truth about the locations of the residents in the home. Once a height sensor based system is deployed in homes, it could potentially be used in our AUTOLABEL approach to cluster events to rooms. Another key challenge is attributing individual fixture usage and activities *inside* rooms to residents identified at entry points in the room; several simple heuristics or solutions from the existing data association literature [50] may be explored to address this problem.

In the future, we envisage an integrated activity recognition system that uses a simple doorway sensor module per room, along with smart meters, to infer the locations and fine-grained activities of residents in the home, along with information about the resources used by each fixture and by each resident. Additionally, emerging single point sensors that can automatically detect the use of low power electronics may be added [57]. Even with such novel sensing solutions, our Bayesian approaches to uniquely identify similar appliances by clustering their power signatures with motion sensor signatures will be necessary and important. Also, in this dissertation, as mentioned earlier, we do not *require* residents to wear tags or carry smartphones at all times. However, in the future, it is possible that residents may use smartphones or other wearable devices some of the time as part of other applications. In such a scenario, it would also be important to develop effective data fusion solutions that leverage both the wearable sensor data with pervasive sensors in the home, to achieve higher accuracy than would be possible from either system alone.

Bibliography

- [1] V.C. Gungor and G.P. Hancke. Industrial wireless sensor networks: Challenges, design principles, and technical approaches. *Industrial Electronics, IEEE Transactions on*, 56(10):4258–4265, 2009.
- [2] T. He, S. Krishnamurthy, L. Luo, T. Yan, L. Gu, R. Stoleru, G. Zhou, Q. Cao, P. Vicaire, J.A. Stankovic, et al. Vigilnet: An integrated sensor network system for energy-efficient surveillance. *ACM Transactions on Sensor Networks (TOSN)*, 2(1):1–38, 2006.
- [3] A. Mainwaring, D. Culler, J. Polastre, R. Szewczyk, and J. Anderson. Wireless sensor networks for habitat monitoring. In *Proceedings of the 1st ACM international workshop on Wireless sensor networks and applications*, pages 88–97. ACM, 2002.
- [4] R.N. Murty, G. Mainland, I. Rose, A.R. Chowdhury, A. Gosain, J. Bers, and M. Welsh. Citysense: An urban-scale wireless sensor network and testbed. In *Technologies for Homeland Security, 2008 IEEE Conference on*, pages 583–588. IEEE, 2008.
- [5] Jiakang Lu, Tamim Sookoor, Vijay Srinivasan, Ge Gao, Brian Holben, John Stankovic, Eric Field, and Kamin Whitehouse. The smart thermostat: using occupancy sensors to save energy in homes. In *Proceedings of Sensys*, 2010.
- [6] A. Wood, G. Virone, T. Doan, Q. Cao, L. Selavo, Y. Wu, L. Fang, Z. He, S. Lin, and J. Stankovic. ALARM-NET: Wireless sensor networks for assisted-living and residential monitoring. *University of Virginia Computer Science Department Technical Report*, 2006.
- [7] Wellaware systems for elderly monitoring. <http://www.wellawaresystems.com>.
- [8] Quietcare systems - living independently. <http://www.quietcaresystems.com>.
- [9] The aware home. Georgia Institute of Technology, <http://www-static.cc.gatech.edu/fce/ahri/>.
- [10] A.P. Glascock and D.M. Kutzik. Behavioral telemedicine: A new approach to the continuous nonintrusive monitoring of activities of daily living. *Telemedicine journal*, 6(1):33–44, 2000.
- [11] X. Jiang, S. Dawson-Haggerty, P. Dutta, and D. Culler. Design and implementation of a high-fidelity ac metering network. In *Proceedings of IPSN*, 2009.
- [12] Y. Kim, T. Schmid, Z.M. Charbiwala, and M.B. Srivastava. Viridiscop: design and implementation of a fine grained power monitoring system for homes. In *Proceedings of Ubicomp*, 2009.

- [13] A. Rowe, M. Berges, and R. Rajkumar. Contactless sensing of appliance state transitions through variations in electromagnetic fields. In *Proceedings of Buildsys*, 2010.
- [14] Y. Kim, T. Schmid, Z.M. Charbiwala, J. Friedman, and M.B. Srivastava. Nawms: nonintrusive autonomous water monitoring system. In *Proceedings of the 6th ACM conference on Embedded network sensor systems*, pages 309–322. ACM, 2008.
- [15] Ge Gao and Kamin Whitehouse. The Self-Programming Thermostat: Optimizing Setback Schedules based on Home Occupancy Patterns. In *First ACM Workshop On Embedded Sensing Systems For Energy-Efficiency In Buildings*, 2009.
- [16] Pike research smart grid deployment tracker.
<http://www.pikeresearch.com/research/smart-grid-deployment-tracker-3q10>.
- [17] Elster smart meter. <http://www.srpnet.com/electric/home/readmeter.aspx>.
- [18] Smart water meters. <http://www.pikeresearch.com/research/smart-water-meters>.
- [19] J. Froehlich, E. Larson, S. Gupta, G. Cohn, M.S. Reynolds, and S.N. Patel. Disaggregated end-use energy sensing for the smart grid. *IEEE Pervasive Computing*, 2010.
- [20] S. Darby. The effectiveness of feedback on energy consumption: A review for defra of the literature on metering, billing and direct displays. *A Review for DEFRA of the Literature on Metering Billing and direct Displays*, 22(April), 2006.
- [21] B. Neenan and J. Robinson. Residential electricity use feedback: A research synthesis and economic framework. *Technical Report*, 2009.
- [22] D.T. Delaney, G.M.P. O’Hare, and A.G. Ruzzelli. Evaluation of energy-efficiency in lighting systems using sensor networks. In *Proceedings of the First ACM Workshop On Embedded Sensing Systems for Energy-Efficiency in Buildings*, pages 61–66. ACM, 2009.
- [23] L.S. Humphries, G. Rasmussen, D.L. Voita, and J.D. Pritchett. Home automation system, April 15 1997. US Patent 5,621,662.
- [24] D.H. Marman and K.B. Liu. Wireless home fire and security alarm system, September 23 2003. US Patent 6,624,750.
- [25] T. Freadman. Hybrid home-entertainment system, February 24 1998. US Patent 5,722,041.
- [26] T. Van Kasteren, A. Noulas, G. Englebienne, and B. Kruse. Accurate activity recognition in a home setting. In *Proceedings of the 10th international conference on Ubiquitous computing*, pages 1–9. ACM, 2008.
- [27] B. Logan, J. Healey, Matthai Philipose, E. Munguia Tapia, and S. Intille. A long-term evaluation of sensing modalities for activity recognition. In *Ubicomp*, 2007.
- [28] E. Munguia Tapia, S. S. Intille, and K. Larson. Activity recognition in the home setting using simple and ubiquitous sensors. In *Proceedings of PERVASIVE*, 2004.
- [29] S.S. Intille, J. Rondoni, C. Kukla, I. Ancona, and L. Bao. A context-aware experience sampling tool. In *CHI’03 extended abstracts on Human factors in computing systems*, pages 972–973. ACM, 2003.

- [30] T.W. Hnat, V. Srinivasan, J. Lu, T.I. Sookoor, R. Dawson, J. Stankovic, and K. Whitehouse. The hitchhikers guide to successful residential sensing deployments. *Sensys*, 2011.
- [31] O. Woodman and R. Harle. Pedestrian localisation for indoor environments. In *Proceedings of the 10th international conference on Ubiquitous computing*, pages 114–123. ACM New York, NY, USA, 2008.
- [32] K. Lorincz and M. Welsh. MoteTrack: a robust, decentralized approach to RF-based location tracking. *Personal and Ubiquitous Computing*, pages 489–503, 2007.
- [33] R. Want, A. Hopper, V. Falcao, and J. Gibbons. The active badge location system. *ACM Transactions on Information Systems (TOIS)*, 1992.
- [34] J.A. Markowitz. Voice biometrics. 2000.
- [35] G. Shakhnarovich et al. Integrated face and gait recognition from multiple views. In *Conference on Computer Vision and Pattern Recognition*. IEEE Computer Society, 2001.
- [36] Predrag Klasnja, Sunny Consolvo, Tanzeem Choudhury, Richard Beckwith, and Jeffrey Hightower. Exploring privacy concerns about personal sensing. In *Proceedings of the Seventh International Conference on Pervasive Computing*, Nara, Japan, May 2009.
- [37] V. Srinivasan, J. Stankovic, and K. Whitehouse. Protecting your daily in-home activity information from a wireless snooping attack. In *Proceedings of the 10th international conference on Ubiquitous computing*, pages 202–211. ACM, 2008.
- [38] V. Srinivasan, J. Stankovic, and K. Whitehouse. Watersense: Water flow disaggregation using motion sensors. 2011.
- [39] V. Srinivasan, J. Stankovic, and K. Whitehouse. Using height sensors for biometric identification in multi-resident homes. *Pervasive Computing*, pages 337–354, 2010.
- [40] United States department of health and human services, HIPAA regulations and standards. <http://www.hhs.gov/ocr/hipaa/>.
- [41] Parks Associates Research and Analysis for Digital Living. Home security system forecasts: 2005 and beyond, November 2005. <http://www.parksassociates.com/research/reports/tocs/2005/security.htm>.
- [42] J. Hall, M. Barbeau, and E. Kranakis. Frequency fingerprinting using phase characteristics of signals. In *IASTED International Conference on Communications and Computer Networks*, July 2003.
- [43] J. Hall, M. Barbeau, and E. Kranakis. Enhancing intrusion detection in wireless networks using radio frequency fingerprinting. In *Communications Internet and Information Technology*, November 2004.
- [44] Kasper Bonne Rasmussen and Srdjan Capkun. Implications of Radio Fingerprinting on the Security of Sensor Networks. Technical Report 536, ETH Zurich IFW, 2006.
- [45] S. Consolvo and M. Walker. Using the experience sampling method to evaluate ubicomp applications. *Pervasive Computing, IEEE*, 2(2):24–31, 2003.

- [46] M. Csikszentmihalyi and R. Larson. Validity and reliability of the experience sampling method. *The experience of psychopathology: Investigating mental disorders in their natural settings*, pages 43–57, 1992.
- [47] D.J. Moore, I.A. Essa, and M.H. Hayes III. Exploiting human actions and object context for recognition tasks. In *Computer Vision, 1999. The Proceedings of the Seventh IEEE International Conference on*, volume 1, pages 80–86. IEEE, 1999.
- [48] P. Turaga, R. Chellappa, V.S. Subrahmanian, and O. Udrea. Machine recognition of human activities: A survey. *Circuits and Systems for Video Technology, IEEE Transactions on*, 18(11):1473–1488, 2008.
- [49] S. Meyer and A. Rakotonirainy. A survey of research on context-aware homes. In *Proceedings of the Australasian information security workshop conference on ACSW frontiers 2003-Volume 21*, pages 159–168. Australian Computer Society, Inc., 2003.
- [50] D.H. Wilson and C. Atkeson. Simultaneous tracking and activity recognition (STAR) using many anonymous, binary sensors. In *The Third International Conference on Pervasive Computing*. Springer, 2005.
- [51] L. Bao and S. Intille. Activity recognition from user-annotated acceleration data. *Pervasive Computing*, pages 1–17, 2004.
- [52] K.P. Fishkin, M. Philipose, and A. Rea. Hands-on rfid: Wireless wearables for detecting use of objects. In *Wearable Computers, 2005. Proceedings. Ninth IEEE International Symposium on*, pages 38–41. IEEE, 2005.
- [53] J.R. Smith, K.P. Fishkin, B. Jiang, A. Mamishev, M. Philipose, A.D. Rea, S. Roy, and K. Sundara-Rajan. RFID-based techniques for human-activity detection. *Communications of the ACM*, 2005.
- [54] Danny Wyatt, Matthai Philipose, and Tanzeem Choudhury. Unsupervised activity recognition using automatically mined common sense. In *AAAI*, 2005.
- [55] J. Froehlich, E. Larson, T. Campbell, C. Haggerty, J. Fogarty, and S.N. Patel. Hydrosense: infrastructure-mediated single-point sensing of whole-home water activity. In *Proc. UbiComp*, volume 9, pages 235–244. Citeseer, 2009.
- [56] J. Froehlich, E. Larson, E. Saba, T. Campbell, L. Atlas, J. Fogarty, and S. Patel. A longitudinal study of pressure sensing to infer real-world water usage events in the home. *Pervasive Computing*, pages 50–69, 2011.
- [57] 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 Ubicomp*, 2010.
- [58] 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. *Pervasive*, 2007.
- [59] G. Cohn, S. Gupta, J. Froehlich, E. Larson, and S. Patel. Gassense: Appliance-level, single-point sensing of gas activity in the home. *Pervasive Computing*, pages 265–282, 2010.

- [60] Toshihiro Takada et al, Proximity mining: Finding proximity using sensor data history. In *WMCSA*, 2003.
- [61] M. Ester et al. A density-based algorithm for discovering clusters in large spatial databases with noise. In *International Conference on Knowledge Discovery and Data Mining*, 1996.
- [62] Water conservation tips, facts and resources. <http://www.wateruseitwisely.com>.
- [63] P.W. Mayer and W.B. DeOreo. *Residential end uses of water*. American Water Works Association, 1999.
- [64] J. Fogarty, C. Au, and S.E. Hudson. Sensing from the basement: a feasibility study of unobtrusive and low-cost home activity recognition. In *Proceedings of the 19th annual ACM symposium on User interface software and technology*, pages 91–100. ACM, 2006.
- [65] X10 home security home automation electronics. <http://www.x10.com>.
- [66] Shenitech water flow meter. <http://www.shenitech.com/>.
- [67] J. Canny. A computational approach to edge detection. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, (6):679–698, 1986.
- [68] H.W. Kuhn. The hungarian method for the assignment problem. *Naval research logistics quarterly*, 2(1-2):83–97, 1955.
- [69] L.J. Heyer, S. Kruglyak, and S. Yooseph. Exploring expression data: identification and analysis of coexpressed genes. *Genome research*, 9(11):1106, 1999.
- [70] Zwave sensors. <http://www.z-wave.com/>.
- [71] G.W. Hart. Nonintrusive appliance load monitoring. *Proceedings of the IEEE*, 80(12):1870–1891, 1992.
- [72] Energy Information Administration. 2005 residential energy consumption survey. <http://www.eia.doe.gov/emeu/recs/contents.html>.
- [73] San Diego Regional Energy Office. Residential Energy Efficiency Measures. [http://energycenter.org/uploads/Residential Efficiency Measures.pdf](http://energycenter.org/uploads/Residential%20Efficiency%20Measures.pdf).
- [74] D. Jung and A. Savvides. Estimating building consumption breakdowns using on/off state sensing and incremental sub-meter deployment. In *Proceedings of Sensys 2010*.
- [75] Ted the energy detective. <http://www.theenergydetective.com>.
- [76] Hamamatsu light diodes. <http://www.hamamatsu.com>.
- [77] J. Polastre, R. Szewczyk, and D. Culler. Telos: enabling ultra-low power wireless research. In *Proceedings of IPSN*, 2005.
- [78] A.D. Smith. Exploring the acceptability of biometrics and fingerprint technologies. *International Journal of Services and Standards*, 2005.
- [79] M.D. Addlesee, A. Jones, F. Livesey, and F. Samaria. The ORL active floor. *IEEE Personal Communications*, 1997.

- [80] J. Jenkins and C. Ellis. Using ground reaction forces from gait analysis: body mass as a weak biometric. *Lecture Notes in Computer Science*, 2007.
- [81] V. Shnayder et al. Sensor networks for medical care. In *Sensys*. ACM, 2005.
- [82] M. Köhler, S.N. Patel, J.W. Summet, E.P. Stuntebeck, and G.D. Abowd. Tracksense: Infrastructure free precise indoor positioning using projected patterns. In *Proceedings of the 5th international conference on Pervasive computing*, pages 334–350. Springer-Verlag, 2007.
- [83] S.N. Patel, K.N. Truong, and G.D. Abowd. Powerline positioning: A practical sub-room-level indoor location system for domestic use. *Lecture Notes in Computer Science*, 4206:441–458, 2006.
- [84] Ben Abdelkader et al. Person identification using automatic height and stride estimation. In *International Conference on Pattern Recognition*, 2002.
- [85] Y. Nishida, S. Murakami, T. Hori, and H. Mizoguchi. Minimally privacy-violative human location sensor by ultrasonic radar embedded on ceiling. In *Sensors, 2004. Proceedings of IEEE*, 2004.
- [86] Jam Jenkins et al. Weakly identifying system for doorway monitoring. Duke Fontiers Poster Session, May 2006.
- [87] Vernier go motion ultrasonic sensor. <http://www.vernier.com/go/gomotion.html>.
- [88] Health and retirement study 2006. <http://hrsonline.isr.umich.edu>.
- [89] A. Perrig, R. Szewczyk, JD Tygar, V. Wen, and D.E. Culler. Spins: Security protocols for sensor networks. *Wireless networks*, 8(5):521–534, 2002.
- [90] L. Eschenaur and V. Gligor. A key-management scheme for distributed sensor networks. In *ACM conference on Computer and communications security*, 2002.
- [91] Hagai Bar El, *Introduction to Side Channel Attacks*, <http://www.hbarel.com/publications.htm>.
- [92] Markus Kuhn. Electromagnetic eavesdropping risks of flat-panel displays. In *Workshop on Privacy Enhancing Technologies*, 2004.
- [93] S. Saponas, J. Lester, C. Hartung, S. Agarwal, and T. Kohno. Devices that tell on you privacy trends in consumer ubiquitous computing. In *Usenix Security Symposium*, 2007.
- [94] J. Hall, M. Barbeau, and E. Kranakis. Detecting rogue devices in bluetooth networks using radio frequency fingerprinting. In *IASTED International Conference on Communications and Computer Networks*, October 2006.
- [95] Jeffrey Pang, Ben Greenstein, Ramakrishna Gummadi, Srinivasan Seshan, and David Wetherall. 802.11 user fingerprinting. In *MobiCom*, 2007.
- [96] Florian Dotzer. Privacy issues in vanet. In *workshop on Privacy Enhanced Technology*, 2005.
- [97] David Chaum. Untraceable electronic mail, return addresses and digital pseudonyms. In *Communications of the ACM*, pages 84–88, 1981.
- [98] P. Kamat, Y. Zhang, W. Trappe, and C. Ozturk. Enhancing source location privacy in sensor network routing. In *Int Conference on Distributed Computing Systems*, 2005.

- [99] Jing Deng Han et al, Countermeasures against traffic analysis attacks in wireless sensor networks. In *SecureComm*, 2005.
- [100] Yi Yang et al, Towards event source unobservability with minimum network traffic in sensor networks. In *WiSec*, 2008.
- [101] Assisted living and residential monitoring network project. University of Virginia ALARMNET project, <http://www.cs.virginia.edu/wsn/medical/>.
- [102] Daniel B. Faria and David R. Cheriton. Detecting identity-based attacks in wireless networks using signalprints. In *Wise*, 2006.
- [103] Anritsu high performance signal analyser. http://www.scs.carleton.ca/jhall2/Publications/anritsu_us.pdf.
- [104] Dark deal hacking wireless video cameras. <http://www.g4tv.com/techtv/vault/features/46880/>.
- [105] R. Lim. Wireless fire sensor network demonstrator. Master's thesis, ETH Zurich, 2006.
- [106] R.F. Dickerson, E.I. Gorlin, and J.A. Stankovic. Empath: a continuous remote emotional health monitoring system for depressive illness. *Wireless Health*, 2011.