

# Room-level Localization of Individuals by Height for Smarthome Applications

---

A Thesis

Presented to

the Faculty of the School of Engineering and Applied Science

University of Virginia

---

In Partial Fulfillment

of the requirements for the Degree

Master's of Science (Computer Science)

by

Erin Griffiths

May 2013



# Abstract

The foundation for many intelligent building applications is the ability to track the movement of occupants for personalized environment conditioning and services. However, today there are no practical tracking solutions that are tolerant to the daily life of an actual home. Most solutions either require occupants to carry wearable devices or perform costly installations that are not practical in a residential home. By monitoring doorways, coarse grained room-level tracking information can be obtained on residents without hindering their daily lives. This thesis presents a doorway tracking system that can track residents in the noisy environment of their own home. Individuals are tracked based on their height, identifying when and in which direction they walk through doorways. The Screen Door System presented in this work outperforms other height based systems in an in-situ environment by 46% and can identify doorway crossings with 93% accuracy.

# Approval Sheet

This thesis is submitted in partial fulfillment of the requirements for the degree of  
Master's of Science (Computer Science)

---

Erin Griffiths

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

---

Kamin Whitehouse, Adviser

---

Mary Lou Soffa

---

Worthy Martin

---

Kevin Sullivan

Accepted for the School of Engineering and Applied Science:

---

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

May 2013

# Contents

<b>Contents</b>	<b>iii</b>
List of Tables . . . . .	v
List of Figures . . . . .	vi
<b>1 Introduction</b>	<b>1</b>
1.1 Overview . . . . .	2
1.1.1 Doorway Sensors . . . . .	3
1.1.2 Data Processing . . . . .	4
<b>2 Related Work</b>	<b>5</b>
2.1 Carried Devices . . . . .	5
2.2 Visual Systems . . . . .	6
2.3 Non-invasive Systems . . . . .	6
2.4 Doorjamb 1.0 . . . . .	7
2.4.1 Ultrasonic Range Finders . . . . .	7
2.4.2 Signal Processing . . . . .	9
2.4.3 Drawbacks . . . . .	12
<b>3 Doorjamb 2.0</b>	<b>15</b>
3.1 Approach . . . . .	15
3.1.1 Crossing Detection . . . . .	15
3.1.2 Noise Filtering . . . . .	16
3.1.3 Height Detection . . . . .	17
3.1.4 Direction Detection . . . . .	18
3.2 Tracking . . . . .	19
3.3 Evaluation . . . . .	19
3.3.1 Experimental Setup . . . . .	19
3.3.2 Results . . . . .	21
3.4 Conclusion . . . . .	22
<b>4 Ground Truth</b>	<b>25</b>
4.1 Radio Frequency Identification . . . . .	25
4.1.1 Consideration For Ground Truth . . . . .	26
4.2 Approach . . . . .	27
4.2.1 False Positive Identification . . . . .	27
4.2.2 True Positive Identification . . . . .	27
4.2.3 False Negative Identification . . . . .	28
4.2.4 Resident Behavior Identification . . . . .	28
4.2.5 Direction Identification . . . . .	29
4.3 Validation . . . . .	29
4.3.1 Experimental Setup . . . . .	30
4.3.2 Results . . . . .	30
4.4 Conclusion . . . . .	32

<b>5</b>	<b>Screen Door System</b>	<b>35</b>
5.1	Infrared Range Finders . . . . .	35
5.2	Approach . . . . .	37
5.2.1	Crossing Detection . . . . .	37
5.2.2	Softening Event Segmentation . . . . .	38
5.2.3	Height Detection . . . . .	38
5.2.4	Direction Detection . . . . .	39
5.3	Evaluation . . . . .	39
5.3.1	Experimental Setup . . . . .	39
5.3.2	Results . . . . .	40
5.4	Conclusion . . . . .	44
<b>6</b>	<b>Smarthome Applications</b>	<b>45</b>
6.1	Case Study: Floor Plan Inference . . . . .	45
6.1.1	Approach . . . . .	46
6.1.2	Evaluation . . . . .	48
6.1.3	Conclusion . . . . .	48
<b>7</b>	<b>Conclusion</b>	<b>49</b>
7.1	Limitations and Future Work . . . . .	49
	<b>Bibliography</b>	<b>51</b>

# List of Tables

3.1 Doorjamb Improvement . . . . .	21
------------------------------------	----



# List of Figures

1.1	Doorway Sensor . . . . .	3
1.2	Data Flow Overview . . . . .	4
2.1	Ultrasonic Range Finders . . . . .	7
2.2	Environmental Noise . . . . .	9
2.3	Baseline Noise . . . . .	10
2.4	Doorjamb Performance . . . . .	12
3.1	Noise Filtering . . . . .	16
3.2	Example Noise . . . . .	18
3.3	Instrumented Doorways . . . . .	19
3.4	Bi-partite Matching . . . . .	20
3.5	Height Accuracy . . . . .	22
3.6	Tracking Accuracy . . . . .	23
4.1	RFID Antenna Placement . . . . .	26
4.2	Proxy Metric Methods . . . . .	28
4.3	In-Situ Experimental Setup . . . . .	30
4.4	Precision and Recall Validation . . . . .	31
4.5	Proxy Metric Validation . . . . .	32
5.1	Infrared Range Finders . . . . .	36
5.2	Infrared Screen . . . . .	36
5.3	Event Detection Accuracy . . . . .	41
5.4	Height Accuracy . . . . .	42
5.5	Direction Accuracy . . . . .	43
6.1	Floor Plan Generation . . . . .	46



# Chapter 1

## Introduction

Buildings of the future are becoming increasingly intelligent. It is now commonplace to have motion controlled lights in all rooms of an office building and to automate a building to shut down at the end of the work day. Motion sensed toilets, faucets, and soap and towel dispensers are becoming increasingly common in new commercial buildings. Behind all of these new innovations is the drive to save energy. The easiest way to save energy is, of course, not to use it. If no people are present in a building to use it, then nearly all of the buildings energy use can be saved by shutting it down. However, people must consume energy in a building to be productive. The problem then becomes how to save wasted energy, when a person does not require this energy to be comfortable or productive.

Determining what energy is wasted energy requires the ability to monitor the activities of people. Motion sensors connected to lights attempt to do this by shutting down when no motion, and presumably no people, are detected in the room. Zoned heating and cooling has been studied in an attempt to perform a similar task with HVAC systems, conditioning a room only if a person is present to need that conditioning. Most of these applications require a room level knowledge of the location of people. In HVAC systems, knowledge of a person's habits, such as they arrive at work at 7:00am, can mean that a room is heated or cooled to their preference before they arrive.

Most of these advancements, however, apply to newer and commercial buildings. According to a recent survey [1], people spend 62% of their time within their own homes. Knowledge of a person's location and activities in a home can provide a foundation for a large subset of current and future intelligent home applications. The same energy saving techniques such as occupancy driven HVAC and lighting control present in many modern commercial buildings, could be performed in homes. Applications such as elderly and patient monitoring, activity recognition, and usage based energy feedback are also possible with knowledge of a

person's location. Practical systems that provide location information of individuals within a residential home are difficult to find today. Motion sensors alone lose accuracy when a person sits still for an extended period of time—a common occurrence in a home. Current state of the art occupancy systems require the monitored individuals to wear some form of tagged identification, often in the form of a powered device such as a RF transceiver, or consist of privacy invasive devices such as cameras or microphones. While these solutions may be acceptable in public or commercial buildings, as people commonly carry identification cards and are monitored by security cameras, this is not true of a home. Requiring residents to carry devices or interact with a monitoring system disrupts their daily lives. Often such devices will be forgotten or residents will refuse to wear them after a time and placing cameras or microphones within a home is seen as an invasion of privacy.

One practical approach to an in-home tracking system that can be considered non-invasive is achieved through doorway sensing. To transition from room to room within home, people must travel through doors. Hence, monitoring devices are only needed in each doorway to achieve room-level knowledge of a person's location. Residents do not need to carry wearable devices and no cameras or microphones are required. By determining the occurrence of a doorway crossing and the identity and direction of the person who made the crossing, individual residents can be tracked as they move throughout their home. This room-level location knowledge can then be used as the foundation for many residential smarthome applications.

Doorjamb 1.0 is a current implementation of this type of doorway sensing system uses ultrasonic range finders to determine when an individual transitions through a doorway, which direction they are going, and their identity [2]. Height is used as a weak biometric to infer the identity of residents. Due to the fact these range finders send signals that can bounce off the objects in the surrounding environment, this system is greatly affected by environmental noise such as doors, people situated near the doorway, and the ultrasonic signals from other doorways. These types of noise are limited in a controlled environment, which is where the height-based doorway sensing system first showed promise. Under the conditions of daily life in home, this noise can occur frequently. The Doorjamb system is greatly affected by this, losing precision as a large number of false crossings are detected. This thesis contributes the following: Doorjamb 2.0, a sensing system based on prior work [2] with optimized processing to account for use in daily life and the Screen Door System, a hybrid system with both ultrasound and infrared sensors.

## 1.1 Overview

Each doorway system presented in this thesis consists of two similar components and builds on the Doorjamb 1.0 system in prior work [2]. The first component is the doorway sensors, enclosures instrumented with

sensors that are placed in the top of each doorway. The second component is the general process that is used to transform raw data into information about residents' locations.

### 1.1.1 Doorway Sensors

The doorway sensors are designed to be snap-in sensors that blend easily into the environment of a home. To encourage use within homes, the doorway sensor enclosure must be compact, able to fit and attach easily to any doorjamb, and relatively invisible in a doorway. Therefore, each enclosure is laser cut to size and folded out of heavy mat board. Once placed in a home, these enclosures can last and continue to function for months.

Height is retained as the weak biometric identifying feature in all doorway sensors systems presented here. While not a uniquely identifying feature in large populations, within most homes there is enough of a height difference to differentiate between individuals [3]. Small changes in height due to shoes or hair styles must also be tolerated when monitoring residents in their daily lives.

The doorway sensing systems in this thesis use only infrared and ultrasonic sensors to track individuals throughout the home. In preparation for future work a variety of other sensors are included in the design of the doorway sensors. First, two PIR motion sensors are included. One of these PIR sensors faces into each room on either side of the doorway. Second, magnetometer sensors are included to determine the global (north, east, south, west) orientation of the doorways. Third, temperature sensors are included for use in future HVAC control. Fourth, light sensors are placed facing into the rooms on either side of the doorway. Most of the enclosures are similar sized, flat, rectangular tubes as shown in Figure 1.1.



Figure 1.1: Doorway sensors are placed at the top of the doorway and designed to blend in. Sensors face down or out towards the rooms and residents.

Each doorway sensors collects and transmits data through the use of a SNAPpy microcontroller. Once data has been collected, it is sent wirelessly to a laptop basestation located within the home and forwarded on to a server at the University of Virginia for analysis.

### 1.1.2 Data Processing

Each system in this thesis goes through a process to transform the raw data collected from the doorway sensors into resident location information. The raw data is first processed into a set of doorway crossings across all doors in a home by a step called signal processing. Signal processing consists of three steps: crossing detection, to detect if a crossing has occurred, height detection, to obtain the height of the resident who crossed, and direction detected, to estimate the resident's direction of travel. The set of doorway crossings is then processed by tracking, which analyzes and converts the crossings into a set of room occupancies made by the residents'. A general overview of this process is presented in Figure 1.2. The signal processing step and its components are the main focus of this thesis. A full description of the tracking step can be found in [2].

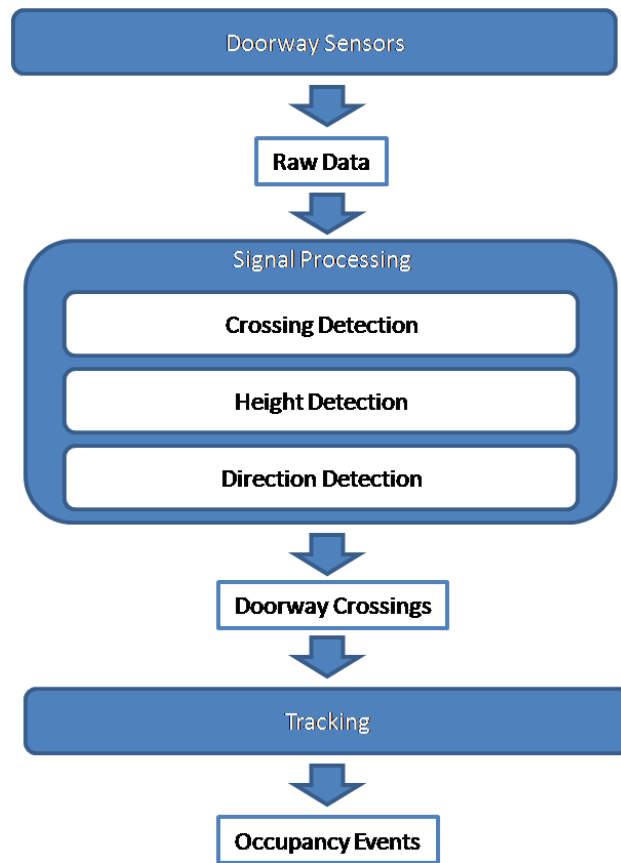


Figure 1.2: Data collected from doorway sensors goes through two steps, signal processing and tracking, to produce information about the residents' locations.

## Chapter 2

# Related Work

Localization of people within a building has been widely studied for a variety of applications. As these applications vary, so does the needed granularity of the localization information. For some applications a very fine grained localization of individuals is needed. For example, if you want to know which of two people in a room are currently using the coffee maker some type of coordinate based localization is needed. Others require less fine grained localization information, such as what room an individual is located in, but not where in that room they are. Examples of these applications include HVAC and room lighting control. Some applications require the ability to assign identities to individuals and some do not.

Due to the large variety of applications that require tracking of people in buildings, a diverse set of localization technologies have been developed. These technologies can be roughly classified based on their interactions with users into three categories: carried devices, visual systems, and non-invasive systems.

### 2.1 Carried Devices

Most systems that track individual movement in buildings today require residents to carry physical devices such as RFID tags, GPS devices, cell phones or battery powered transceivers. The Active badges system [4] localizes people or objects using small infrared (IR) transmitters worn by people or placed on objects. These signals are collected by receivers located in the area and the estimated distance from the transmitters is used to localize their location. RADAR [5], Motetrack [6], and ActiveBats [7] perform similar transmitter/receiver based localization using radio transceivers. Cricket [8] uses ultrasound and wall- and ceiling-mounted beacons to transmit location based information to the wearable devices that use the information to localize themselves. Other systems use the varying strength of transmission power or signal strength to localize RFID tags [9, 10] or wearable readers [11] in space. Some studies look at unmodified carried devices, such as the smartphones

in [12] where normal, periodic Wi-Fi transmissions are detected by stationary basestations, to perform localization.

In a commercial environment these tracking solutions can be highly practical. Many people are already required to carry some form of identification at work and often carry their own cell phones. Some of these systems have also been shown to obtain fine grained localization information that identifies where within a room an individual is located [13]. However, within homes these wearable devices are often not feasible [14]. Residents may need or want to remove the devices for activities such as showering, changing clothes, or sleeping. They might be forgotten, require battery changes, or, if the device is a cell phone, left in a charging station while the residents perform other activities. In these cases wearable device systems are forced to monitor the devices' location only, which may be very different from that of the person they are trying to monitor. Hours or days of tracking information can be lost in this situation, preventing intelligent applications such as environmental conditioning or medical monitoring from operating successfully.

## 2.2 Visual Systems

Visual tracking systems are installed within the environment surrounding occupants and collected information-rich data about what they can see. People being tracked do not need wear or carry any sort of device as their location is determined from this information-rich data alone. Some of these systems rely on facial and/or gait analysis from recorded video [15, 16] to determine individual identity. In one visual system, ceiling mounted sensors can be used to track the movements of people in a single room [17]. However, most video recording devices systems require light to function and all require line-of-sight. In a home environment, many cameras would need to be placed in the rooms of a home to produce accurate tracking and though thermographic cameras can be used at night, they are very expensive. Additionally, while video recording can be common and tolerated in public spaces, it is often considered a privacy invasion in one's own home, especially when there is the possibility that the data could be hacked.

## 2.3 Non-invasive Systems

Like visual systems, non-invasive tracking systems are installed in the environment where tracking needs to occur and do not require wearable devices. They differ from video devices due to the lack of information-rich recording. Therefore, alternate and weaker biometrics must be used. Non-invasive smart floor systems identify individuals by weight [18, 19], but are often too expensive for residential use due to significant installation

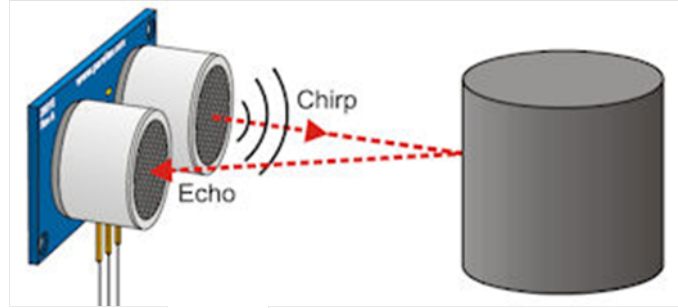


Figure 2.1: Ultrasonic range finders detect the distance to an object by measuring the return time of a reflected signal.

costs. Also, the weight of an individual often fluctuates due to daily eating and exercise habits or carried objects.

Many non-invasive systems track people without inferring identity. Motion sensors can be used to localize the presence of people in a single room or find the location of a single person within a room [20]. The attenuation of RF signals caused by the human body can also be measured to localize a person within a single room [21]. The main drawback to these systems is that they cannot detect the identity of any individual being tracked. Additionally, they can often not detect if there is a single occupant or many.

## 2.4 Doorjamb 1.0

Doorjamb 1.0 [2] falls into the non-invasive category. Height, as the weak biometric used for identification, does not have the major fluctuation issues found in weight. Additionally, as all sensors are placed within doorways and collect only distance measurement data, no person is required to carry a tracking device or endure privacy invasive visual recording. Prior work on a doorway sensor system uses only ultrasonic range finders to detect crossing events, the height of the person responsible, and their direction.

### 2.4.1 Ultrasonic Range Finders

Ultrasonic range finders measure the distance from the closest object by way of an acoustic pulse. These devices generally consist of a transmitter and receiver pair. The transmitter sends out an acoustic pulse by applying a high voltage to a metal diaphragm. The device measures the time it takes for the signal to return to the receiver, which is then converted into a measure of distance. The PING))) Ultrasonic Distance Sensors deployed here achieve an accuracy within 1cm of the actual distance to the measured object.

Each acoustic signal transmits as a wide beam, forming a cone of sensed space as shown in Figure 2.2. When placed in a doorway facing downwards, this sensor can measure the height of a person present in the

doorway given the height of the doorway is known. However, due to the conical shape of the sensing region there are a number of errors that can occur when attempting to measure a person's height in a doorway. First, a geometric error can occur, where the person is not standing directly between the ultrasonic sensor and the floor. When this occurs, the person's height is still measured from the top of their head, but the distance from the sensor becomes longer due to the horizontal displacement. Therefore, the person's height measurement is slightly shorter than it should be. Second, a multi-path error is possible when two or more ultrasonic sensors are placed in the same doorway. Since the sensors fire in succession, signals from one sensor may bounce off the person's head or shoulders and be received by another sensor in the doorway. Because this received signal is not a signal sent out by the second sensor, it can be picked up immediately by the second sensor after it sends its own pulse. In this case, the return time may be much shorter than an actual signal sent out by the sensor would be, making the measured height much taller than the person. Third, sample rate errors can occur, where a person was standing directly under the sensor, but was not sampled exactly at that time. In this case an error occurs for the same reason as a geometric error, but instead of horizontal displacement parallel to the doorway, this displacement is perpendicular to the doorway. Each of these errors can cause a person's height estimate to be larger or smaller than the actual height of the person. This issue can be further exacerbated by a person's walking pattern, i.e. hunching or having a wide, dipping stride.

Ultrasonic range finders can also detect objects that are outside of the doorway, as shown in Figure 2.2. Objects such as doors or nearby furniture can be detected and measured by the doorway sensors. People who walk by a doorway, but not through, can also be detected and measured by the sensors. Additionally, ultrasonic signals from nearby sensors can be detected as measurements when the signal bounces off a surface like the floor. Many of these measurements can appear within the human height range and, especially for a person walking by, can produce a pattern of height measurements that are indistinguishable from an actual crossing.

The conical shape of the ultrasound sensors does provide information that can be used to determine a person's direction of travel. By tilting the cone to one side of the doorway, a person can be detected before they cross the threshold. Due to the geometric properties of the cone, each measurement of the approaching person increases in size as they come closer to the doorway. Once the person passes through, their height is no longer detected by the sensor as there are no signals being sent out on that side of the doorway. Detecting if this progression is increasing or decreasing over time indicates if the person entered the doorway through the side the ultrasonic cone is facing, and therefore which direction they traveled through the doorway.

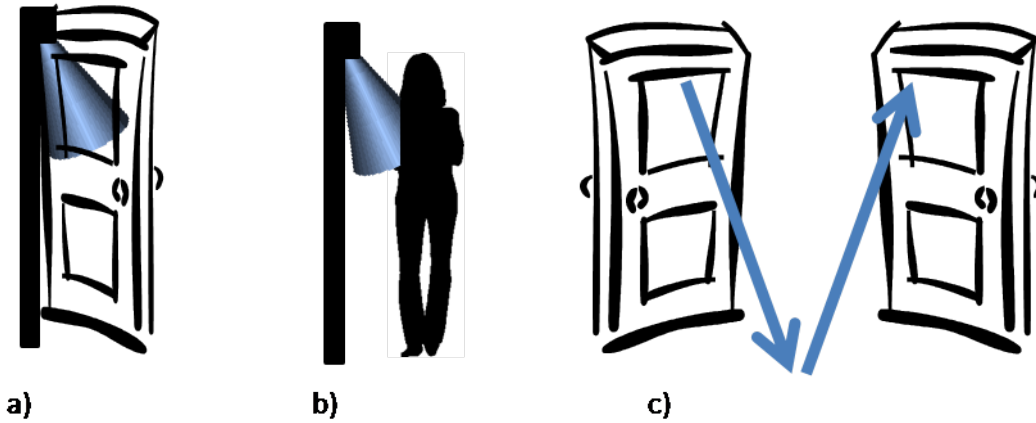


Figure 2.2: Ultrasonic sensors experience noise from the environment such as (a) doors attached to the doorway, (b) people who walk by but not through, and (c) signal reflections from other sensors.

### 2.4.2 Signal Processing

The Doorjamb 1.0 doorway sensor created its signal processing algorithms in a controlled environment. The goal of these signal processing algorithms is to take the raw data produced by the ultrasound sensors and detect when a person walks under a doorway. These doorway crossings can be described as a set events  $E$  where  $e_i \in E = (door_i, time_i, height_i, direction_i)$ : a 4-tuple with information about what door the crossing occurred at, what time the crossing occurred at, and the estimated height and direction of the person traveling through the door.

#### Noise Filtering

Signal processing for the controlled environment based ultrasound-only system relies on a few key insights. The first is that most readings detected by the ultrasonic sensors are baseline noise. Baseline noise is measured when the ultrasonic sensor detects a static surrounding environment. Commonly, this static environment is the floor beneath the doorway or the lower part of the doorjamb that continuously gets observed by the sensors when no person is nearby. A person walking through the doorway is measured with a height much higher than the baseline noise as shown in Figure 2.3. The second insight is that this noise is far more common than height readings from people. If a person crosses a door 1000 times in a day and takes only 1 second to do so then a doorway sensor should measure a value belonging to a person during about only 1 percent of a single day.

Since the doorway system is only interested in the values obtain for people, baseline noise should be eliminated from consideration for events. This is done by creating an inverse mapping between the frequency of occurrence of a measurement and its weight. A normalized kernel density function

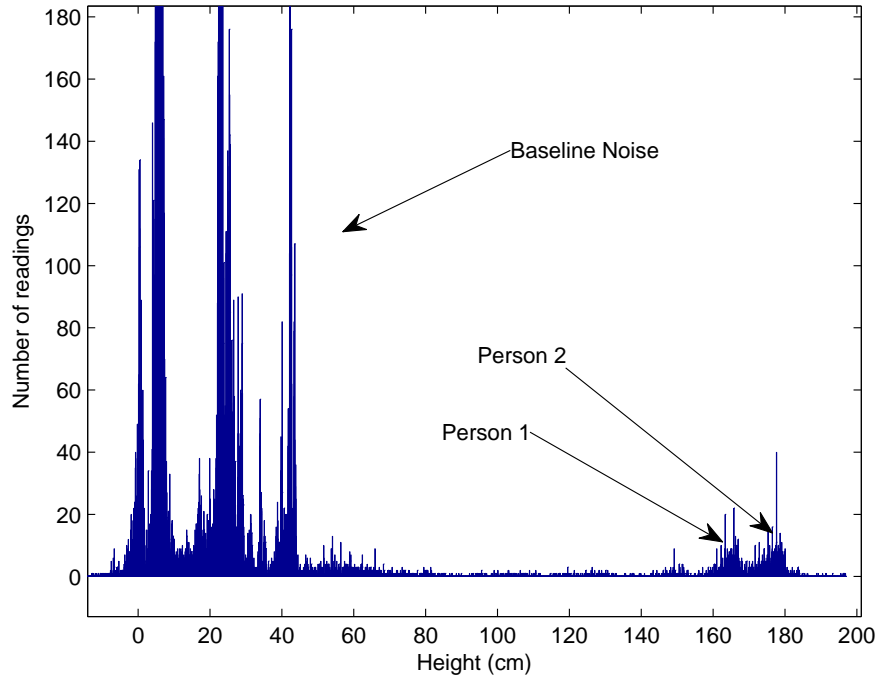


Figure 2.3: Baseline noise can occur at different levels from reflections off surrounding objects and the floor. Height measurements are closer to the sensors then this noise.

$$f_h(x) = \frac{1}{nh} \sum_{i=1}^n K\left(\frac{x - x_i}{h}\right) \quad (2.1)$$

estimates the probability density function of some random input. The parameter  $n$  is the total number of points in the dataset,  $K$  defines the specific kernel function, and  $x_i$  is the set of data where  $i$  is the measurement value. This value is then inverted to weight the less common measurements more. To ensure that uncommon measurements are weighted much higher than baseline measurement, the difference between these values is exaggerated with an exponential-log function

$$weight(x) = (\log(1 - f_h(x)) + 1)^p \quad (2.2)$$

where  $x$  is a particular measurement and the logarithm is raised to the power  $p$ . The value for  $p$  was determined experimentally and set at 1,500.

### Crossing Detection

Doorway crossings are detected by filtering out all uninteresting measurement values with a low weight and examining the remaining data for patterns indicative of a doorway crossing. First, all points weighted

below .1, determined experimentally, are removed. Second, the remaining points are clustered into temporal groups. These values are clustered using DBSCAN, a common algorithm for density-based spacial clustering of applications with noise. This algorithm has two parameters, the minimum points in a cluster and the minimum radial distance to neighboring points in a cluster. These parameters were set experimentally to 2 and 0.1 seconds respectively. Once the measurement points have been clustered into temporal groups, events are identified within these groups by a second run of DBSCAN with a 0.2 second minimum distance. These detected events are then passed on to the height and direction detection algorithms.

### Height Detection

Height is extracted from these events by clustering the data by measurement value. DBSCAN is used to cluster these values with a radial distance of 10cm. Clusters that appear at waist level and below are removed from consideration for height. Clusters that appear above the known maximum height of the residents are also removed to filter any multi-path errors from adjacent sensor signals. The remaining cluster is chosen to represent the true height cluster of the event. The maximum measurement within that cluster becomes the height,  $height_i$ , of the event. The median time of the height cluster becomes the time  $time_i$  of the event.

### Direction Detection

Direction is detected by leveraging the tilt of the ultrasonic sensors. Several features from the data produced by tilted ultrasonic sensors can indicate direction, including maximum height relative to the median, the long-tail direction, and the slope of the regression line. Each of the three feature computations are computed independently and then vote on the direction of travel. For example, a negative slope in the regression line indicates that a person has walked away from the ultrasonic sensor, in the direction that the ultrasound is tilted. As the person's maximum height was obtained when they first entered the ultrasound's range directly under the doorway, later measurements in time are of them walking through and leaving the ultrasound range.

There are three states of direction determined by these algorithms represented by 1,-1, and 0, or IN, OUT, and UNKNOWN, respectively. The vote from each direction algorithm is then summed to form an estimated direction  $direction_i$  where  $-3 \leq direction_i \leq 3$ . Later tracking algorithms weight direction estimates with more extreme values more highly, as all direction metrics have to agree to produce those numbers.

### Stream Fusion

This original signal processing algorithm processes each ultrasonic sensor independently to detect doorway crossings. Stream fusion merges the doorway crossings detected by each sensor in a single doorway. Events

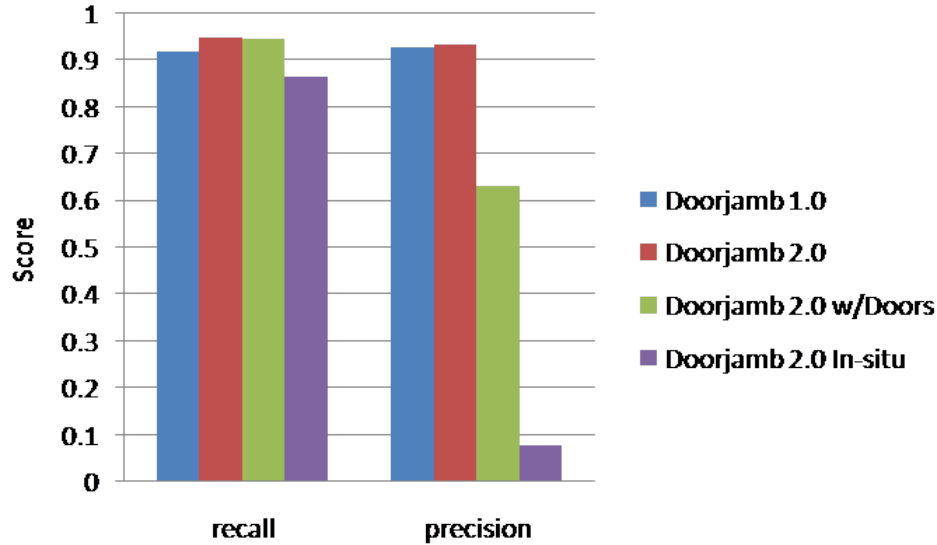


Figure 2.4: Even with a new processing algorithm to account for an in-situ environment, the precision of the Doorjamb system suffers with the introduction of moving doors and a full in-situ setting.

from one stream that have times  $time_i$  within 1 second of an event in another stream are merged. Merged events take the maximum height of all merged events. All un-merged events are included unchanged in the final event stream. These fused events now form the set of doorway crossings  $E$  where  $door_i$  for each event  $i$  is the door containing all ultrasonic sensors that were fused to form the event.

### 2.4.3 Drawbacks

Designing a doorway system in a controlled environment can cause a number of unexpected problems when the system is moved to an in-situ environment. Because of the conical sensing region of the ultrasonic range finders environmental noise such as doors and nearby furniture can be misinterpreted as a crossing event. Additionally, in a real home environment signal processing cannot expect the baseline noise to always remain the same and at a lower value than the heights of people. The largest example of this is that of doors that are attached to doorways. When a door is closed or partially closed the value measured by the ultrasonic sensors can be at or above the height of any residents within the home. Since doors can be left closed for days or more, any function weighting noise based on the inverse of its density of occurrence would remove those values as noise. Since doors can be measured at heights similar to people, all measurements of people at those values would be weighted as noise as well and removed.

This thesis looks to improve the results of Doorjamb 1.0 by first modifying the processing of the Doorjamb system to create Doorjamb 2.0. As you can see from Figure 2.4 even with the new modified processing

algorithm in-situ noise significantly affects accuracy results. The inclusion of doors and other in-situ noise greatly decreases the accuracy of a Doorjamb doorway system. Because of this, this thesis also explores a hardware addition to the doorway sensing system, the Screen Door System, to mitigate these issues.



## Chapter 3

# Doorjamb 2.0

Doorjamb 1.0 used ultrasound sensors in isolation to perform resident tracking as described in Section 2.4. As this system was first tested in a controlled environment, the signal processing algorithm did not take into account environmental noise such as closed doors. To ensure that the addition of infrared sensors to the doorway sensor system was not wasteful, the signal processing of the Doorjamb system was first redesigned to handle an in-situ environment.

### 3.1 Approach

In this new signal processing algorithm all ultrasound readings from all sensors on a single doorway are fused before processing begins. This fusion is based on the insight that usually multiple ultrasonic sensors will detect the same person during a crossing event. Though only one of the ultrasonic sensors should detect the height of the person, the surrounding sensors can provide supportive information for the event, such as measurements of the person's shoulder or arms. The final output of this signal processing algorithm is still the event set  $E$  where  $e_i \in E = (door_i, time_i, height_i, direction_i)$ : a 4-tuple with information about what door the crossing occurred at, what time the crossing occurred at, and the estimated height and direction of the person traveling through the door.

#### 3.1.1 Crossing Detection

The algorithms created for Doorjamb 2.0 seek to detect doorway crossings without relying on any information about the baseline noise. The key insight that allows these algorithms to function is that all measurement values that are caused by people should occur within the adult human height range, 137-198 cm (4 ft, 6 in - 6 ft, 6 in). Though this insight eliminates the detection of children, who are too short, it means that

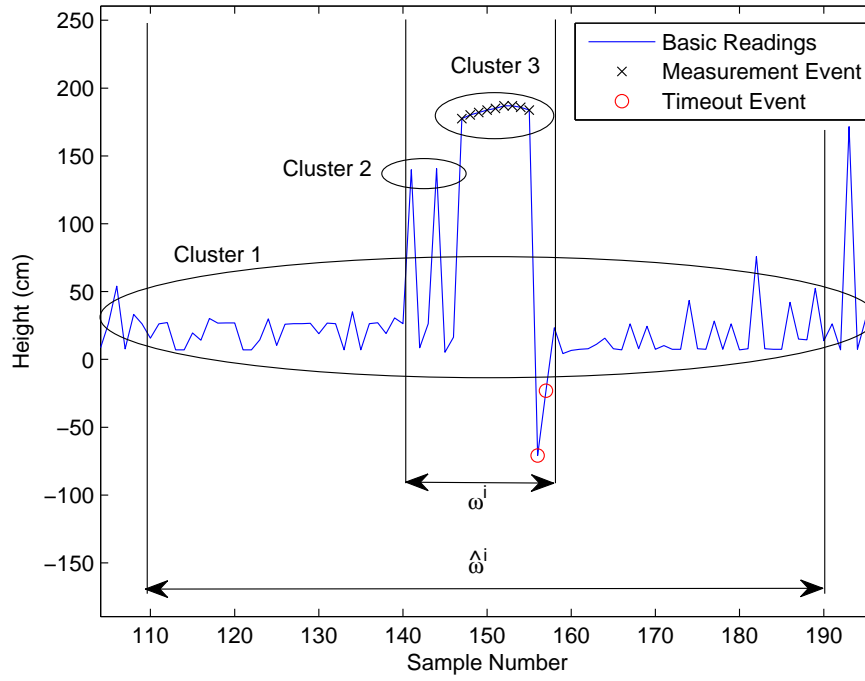


Figure 3.1: The noise filtering algorithm removes baseline values from window  $\omega^i$  if they cluster with values from window  $\hat{\omega}^i$ .

no extended period of data needs to be looked at and weighted to find interesting height measurements. Timeout measurements, where a person's body blocks the signal and it never returns to the sensor, producing a negative height value, are also included in this set of interesting height measurements. An example of such values is depicted by  $x$ 's and  $o$ 's in Figure 3.1.

These interesting measurements are detected and then clustered temporally using DBSCAN, where the minimum distance between any two points is  $400msec$ . A time window  $\omega_i$  is defined around each cluster using the range  $\omega_i = [\min(c_i) - 200msec, \max(c_i) + 200msec]$ , and a new event is defined to be all sensor data from that range. A window  $\omega_i$  is illustrated by the inner bars in Figure 3.1: the event includes all values within that window.

### 3.1.2 Noise Filtering

Ultrasonic range finders can be very noisy and a single outlier value will cause the algorithm above to produce a false detection. To prevent the creation of false doorway crossings from these clusters, filtering is performed to ensure that every detection has a least two values substantially different from the baseline measurements. Baseline measurements are identified based on the  $200msec$  region of data immediately before and after the

detection event:  $\hat{\omega}_i = [\min(e_i) - 200msec, \max(e_i) + 200msec]$ , where  $e_i$  represents all the clustered data in  $\omega$ . Any values that are similar to these baseline regions are removed from the potential doorway crossing. This processes removes both commonly found baseline data, such as that measured at floor level, and baseline data that may change throughout collection, such as a door being open or closed. This  $400msec$  time window will not merge events from two consecutive doorway crossings unless people are walking unusually close to each other.

A window  $\hat{\omega}_i$  is illustrated by the outer bars in Figure 3.1. All sensor values in  $\hat{\omega}_i$  are clustered based on height using DBSCAN and any values that are clustered together with values in the region  $\hat{\omega}_i - \omega_i$  are removed. In the example shown in Figure 3.1, the data in window  $\hat{\omega}_i$  produces three height clusters: 1, 2, and 3. Only cluster 1 includes values in window region  $\hat{\omega}_i - \omega_i$ . Therefore, any values in cluster 1 are removed from the possible event as baseline noise. If less than two values remain in  $\omega_i$  after this filtering, the event is removed.

In addition to random noise, the ultrasonic range finders are also subject to periodic noise due to reflections from nearby environmental features, such as door trim, a hinge, or a nearby shelf. An example of such noise is shown in Figure 3.2, where one doorway sensor frequently picked up reflections from a nearby object with measurements similar to a human height. The key insight about these reflections is that they are consistently the same value and never negative. Therefore, by creating a third window that extends 30 seconds on either side of the potential event, these periodic measurements can be identified and removed. Any measurement in the potential event that is positive and identical to a measurement in this one minute window is removed from the event. People crossing a doorway typically produce a large range of noisy values that are rarely exactly the same as another event. Therefore, even if multiple doorway crossings are contained in this third window the filtering process will not remove most measurements that are not noise. Again if less than two values remain, the event is removed.

Finally, the noise from most closed doors can be eliminated simply by removing any events where the event duration is longer than 6 seconds. This filter may eliminate some crossing events where a person stands in the doorway for an extended period of time before crossing, but it is assumed that this type of behavior is relatively infrequent. After all filtering has been performed each remaining event's median time becomes the time of the doorway crossing  $time_i$  at door  $door_i$ .

### 3.1.3 Height Detection

Once the detection events have been created, a height is estimated for the person responsible for that event. This process is similar to the Doorjamb 1.0 height detection in Section 2.4.2. All height values in an event

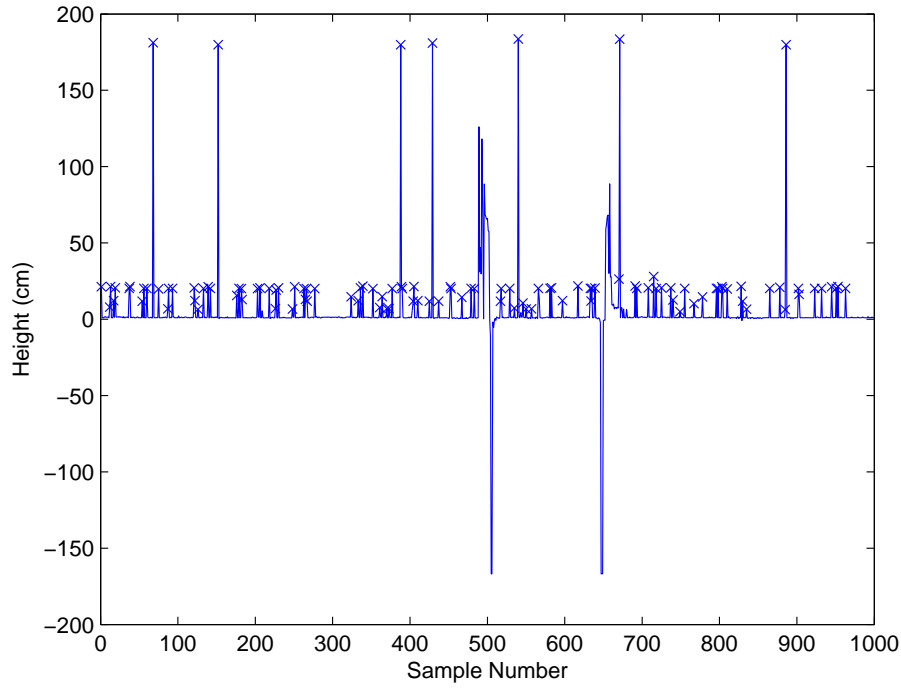


Figure 3.2: Nearby objects reflect acoustic energy and cause random noise at about 180 cm, similar to the height of one of the test subjects.

are clustered with DBSCAN and the maximum height for the maximum cluster is chosen as the event height,  $height_i$ . This process filters out any multi-path errors, where a height taller than that of the actual person is observed. However, if the person walks very quickly underneath the sensors, and their height is recorded as an outlier, they will be assigned a shorter height for this event. If no maximum cluster is formed because of sparse measurements of the person's height during the event, the maximum outlier event within the human height range is chosen as  $height_i$ .

### 3.1.4 Direction Detection

The final feature to be extracted is direction. The main algorithm of this component remains unchanged from that of the direction detection of the Doorjamb 1.0 algorithm in Section 2.4.2. Three different but related algorithms vote on the direction of the crossing and the produce a direction  $direction_i$  where  $-3 \leq direction_i \leq 3$ . Only one thing has changed: values below 140cm are not used to calculate direction. This prevents human movement, like swinging limbs or someone carrying a phone, from interfering with the calculation of direction. After direction has been obtained, the crossing event  $e_i = (door_i, time_i, height_i, direction_i)$  joins the set of detected crossing events  $E$ .

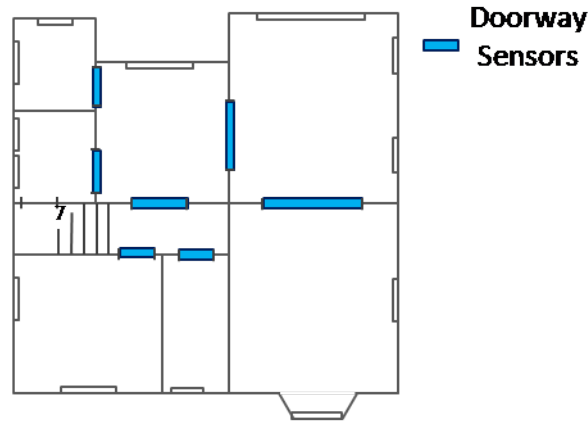


Figure 3.3: Seven internal doorways were instrumented for this experiment.

## 3.2 Tracking

To remove ambiguities from the results both signal processing outputs from the old and new Doorjamb systems were fed through a tracking algorithm. The ambiguities were caused by a variety of errors produced by the signal processing algorithms. These errors included false positives, where the signal processing algorithm detected a crossing when no crossing was made, false negatives, when a crossing that did occur was not detected, and various height and direction errors. The goal of this algorithm is to remove false positives, tolerate the presence of false negatives, and account for errors in height and direction while tracking. Tracking then creates a set of room occupancies recording where a resident was located over the course of the experiment. This algorithm was run unchanged on the signal processing results from both Doorjamb algorithms. A full description and analysis of tracking can be found in [2].

## 3.3 Evaluation

Both signal process algorithms for the Doorjamb system were evaluated in a controlled environment. The final evaluation looked at the output of tracking, where a person has been assigned to a doorway crossing event based on identity information from height and the path constraints of the home.

### 3.3.1 Experimental Setup

Seven doorways within a single home are utilized for this experiment. Data was collected at each doorway from doorway sensors designed for each specific doorway. Most doorway sensors contained two ultrasonic sensors and were attached to a flat doorjamb parallel to the floor. One doorway was wider than the others

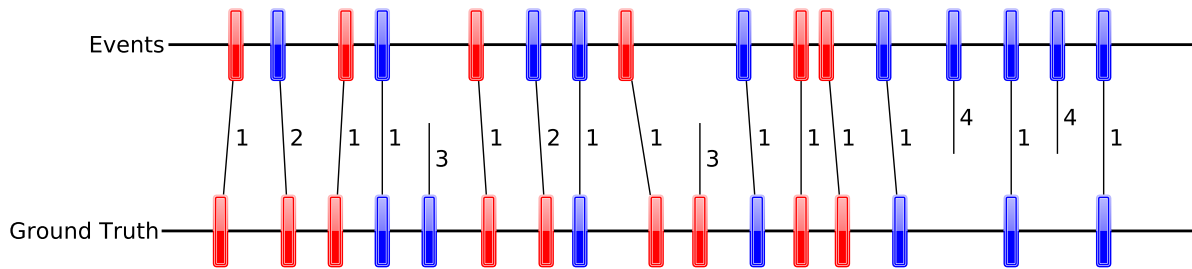


Figure 3.4: Bi-partite matching is used to address timestamp mis-matches between observed events and ground truth [2].

and contained three ultrasonic sensors. Another doorway was a wide, high arch and was fitted with a curved enclosure containing two ultrasonic sensors. All doorway sensors logged their data to a basestation located at the experimental home. This data was then forwarded to a server at the University of Virginia for analysis.

Three controlled experiments were run collecting data from all doorways over a period of one hour. Three people of different heights (180cm, 166cm, and 160cm) participated in the experiments. For each experiment a different pair of the three individuals walked in a random pattern throughout the home, recording 500 doorway crossings each. In total, 3,000 doorway crossing were recorded. Doors were left open and untouched for the entirety of the experiment.

Ground truth was collected manually by the participants as they traveled through the home with hand-held touch-screen phones. A custom build app displayed the layout of the home. Whenever a participant walked from one room to another, they would click on the room they were entering and the crossing would be recorded with the exact identity, direction, and time. Occasionally an incorrect room would be touched by accident. On these occasions participants would immediately touch the correct room before continuing. Because these mis-touches do not follow the participant's path and appear as if the person had teleported, the 20 that were found were easily cleaned from ground truth.

As ground truth was collected by a different system than the doorway data—and by humans—the times that each event occurs is not a direct match between the two systems. In order to compare the doorway results to ground truth, bi-partite matching is used to match doorway crossings to their corresponding ground truth events. Figure 3.4 shows a sample of this matching. All matches labeled 1 show a successful match between the empirical data and ground truth where the doorway, time, and responsible person match. All matches labeled 2 show a successful match but for the wrong person assignment by the doorway system. Labels 3 and 4 show false negatives and false positives, where the doorway system failed to detect an event or detected one that didn't occur.

Signal Processing Algorithm	Precision	Recall
Doorjamb 1.0	92.9%	91.9%
Doorjamb 2.0	93.6%	94.8%

Table 3.1: Precision and recall improved 2-3% over the original Doorjamb system.

### 3.3.2 Results

Doorjamb 2.0's signal processing improves the results of the experiments in two ways. First, small improvements are observed throughout the components of the signal processing step. Second, a significant performance increase was found when the two systems were compared after the results were fed through tracking.

#### Signal Processing

The doorway sensing systems are evaluated using two complementary metrics: precision and recall. Precision represents the number of false positives found by the system, where a crossing was detected but not actually made, and is calculated from the intersection of the ground truth events, and the detected events divided by the number of detected events.

$$precision = \frac{\text{detected events} \cap \text{ground truth events}}{\text{detected events}} \quad (3.1)$$

Recall is calculated from the intersection of the ground truth events and the detected events divided by the number of ground truth events. It represents the number of false negatives found by the system, where a crossing was made, but not detected.

$$recall = \frac{\text{detected events} \cap \text{ground truth events}}{\text{ground truth events}} \quad (3.2)$$

Both metrics showed improvement with Doorjamb 2.0's ultrasound signal processing algorithm as seen in table 3.1. Filtering noise only when it is the baseline around that specific event and removing noise caused by periodic detections of the environment, both necessary pieces for an in-situ deployment, improve precision and recall even in a non-in-situ environment.

To analyze the accuracy of the height extraction the height value of each matched event is assigned to the person who created the corresponding ground truth event. Histograms of the heights for each system are shown in Figure 3.5 with the results from Doorjamb 1.0 on the left and the Doorjamb 2.0 on the right. All participants' heights are grouped better around their respective heights in Doorjamb 2.0, however a large area of overlap still exists, especially in the two shorter heights. This overlap will cause the person responsible

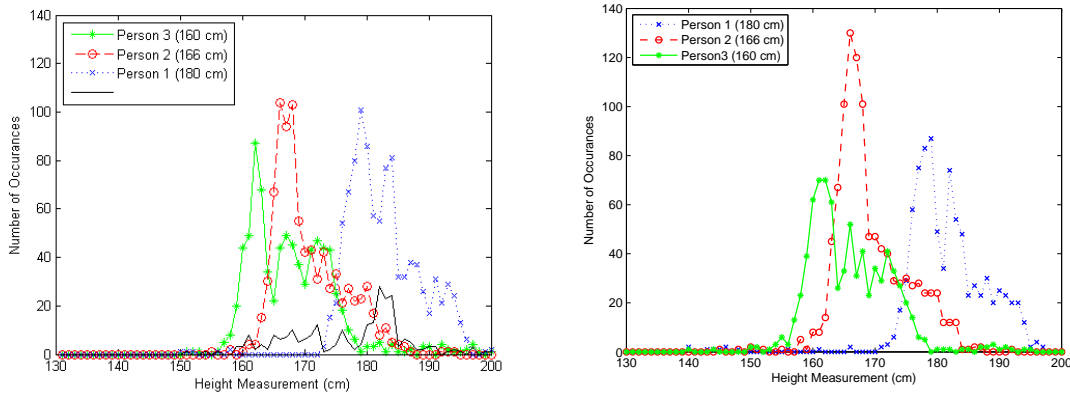


Figure 3.5: The new height detection (right) forms tighter groups around the participants' heights [2]. The overlap seen in the old height detection algorithm (left) still occurs.

for a doorway crossing to be ambiguous and is hence why tracking is necessary to assign an accurate room location for each individual.

Direction accuracy depends greatly on the accuracy of height estimation. If a height is unable to be determined, then no direction is able to be estimated. If the final height is chosen incorrectly, then the direction algorithms may lose accuracy due to incorrect estimates from one or more of the three algorithms. Of the direction estimates that can be determined, Doorjamb 1.0 estimates the direction with 70% accuracy. Doorjamb 2.0 estimates the direction with 81% accuracy. This is likely due to the improved accuracy of the height measurements that direction depends on.

## Tracking

All of the small improvements to the signal processing results produce a large improvement in the tracking results. Figure 3.6 shows the results for both signal processing algorithms after undergoing tracking. The accuracy here represents the number of transitions that were correctly assigned to a participant of the experiment. The precision, recall, height errors, and direction errors all affect the end results produced by tracking. By improving each of these components slightly, tracking is significantly better able to disambiguate events and assign them to the correct individual.

## 3.4 Conclusion

Taking an in-situ environment into account increased the tracking accuracy of residents from an average 75% to 90%. However, these results were produced in a controlled environment. Normal noise from daily life was not included in this experiment: no doors were closed, no participant hovered in the doorway, no person

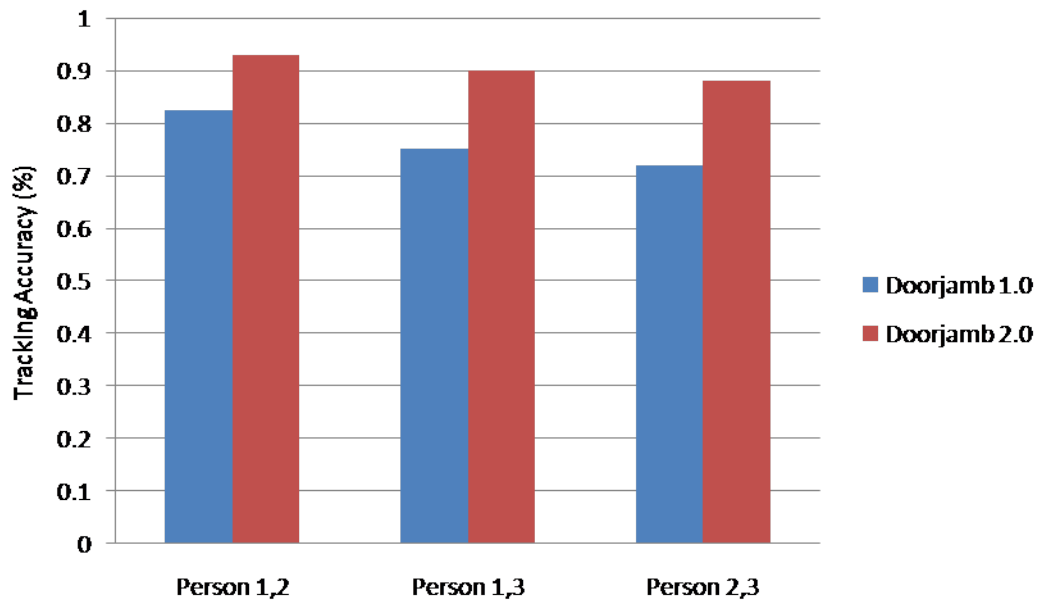


Figure 3.6: Tracking accuracy improves about 20% with Doorjamb 2.0's signal processing.

stood or walked in front of a doorway without going through. Each of these in-situ issues has the potential to significantly reduce the accuracy of this Doorjamb system. Data from an in-situ environment must be collected and analyzed to determine if tracking residents using this hardware can be done in the daily life of residential home.



## Chapter 4

# Ground Truth

In order to evaluate any system scientifically you need ground truth. Ground truth is a huge obstacle to performing in-situ studies with humans. The very act of collecting ground truth during the study may change the results of the experiment as people’s behavior changes to accommodate ground truth collection. Previous studies evaluating the activities or location of people in homes have depended on ground truth methods like manual logging or video recording. Both options can change the behavior of the observed individuals and prevent the collection of true in-situ data. This thesis is not immune to these problems and does not solve them. However, this section describes the proxy metric that has been created to approximate an evaluation with ground truth. The only imposition placed on the residents of the experimental home are bracelets containing radio frequency identification tags that must be worn around each ankle, thus minimizing the affect of ground truth collection on data collection.

### 4.1 Radio Frequency Identification

Radio frequency identification (RFID) wirelessly detects and tracks tags through the use of radio-frequency electromagnetic fields. As used here, when a tag comes within the sensing range of an RFID antenna it transmits its identification information to the antenna. In this way tags can be tracked as they move throughout an area of space monitored by RFID antennas.

Mid-range RFID antennas placed below the floor on either side of every doorway within a home, as shown in Figure 4.1, can identify when a tag is near a specific doorway. These tags are contained within bracelets worn around the ankles of each participant of the in-situ study. The bracelets require no batteries and are waterproofed so they can easily be worn for an extended period of time. Each person is uniquely identified by the tags they are wearing.

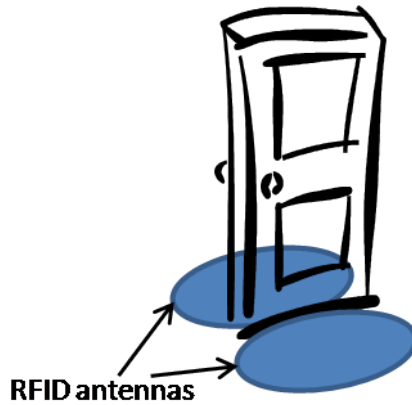


Figure 4.1: RFID antennas placed on either side of the doorway collect identity information about the residents who cross.

#### 4.1.1 Consideration For Ground Truth

Ideally, when a resident crosses over a doorway the antennas on either side of the doorway should detect the tags in their bracelet. Depending on which side was identified first, the direction of travel could be determined and identity could be obtained from the tags. However, testing showed that the current setup of the RFID system made it infeasible for use as ground truth. Despite this, RFID could still provide useful information. A full description and analysis of the RFID based tracking system can be found in [22]. This thesis leverages the infrastructure of this work, but does not use any of the algorithms or processing.

##### **Con: Sensing Region**

The signal strength of RFID tags changes when placed close to the human body. When an RFID tag is within millimeters of human skin the signal is often attenuated, causing the effective sensing region of the RFID antenna to shrink. If the bracelet is worn loosely or is further separated from the skin by pants or socks, the sensing region on the antenna can increase and overlap with the other side of the doorway. This changing sensing region can mean that gaps of un-monitored space can appear within a doorway. Typically, a tag is detected by at least one of the antennas surrounding the doorway, but may not be detected by an antenna on the opposite side. This can cause a crossing to go undetected or make it impossible to determine direction if the crossing is detected.

Because of these issues the RFID system required its own set of signal processing algorithms and evaluations. False positives, when a person walked by the doorway and was detected by the sensing region on both sides of the doorway, had to be balanced with false negatives when someone did walk through the doorway but went undetected by one region. These issues proved RFID to be too inaccurate for use as ground truth.

**Pro: Proximity Identification**

Though the RFID system couldn't produce a ground truth capable set of doorway crossing to compare to the doorway systems, it could provide useful information about the movements of people for our proxy metric. Any time a tag is detected by an antenna, the person wearing that tag must be in proximity to that antenna, and therefore that doorway. When no tags are detected, then no resident is near that doorway. Also, the tags provide accurate information on the identity of the resident who is in the proximity of the antenna.

## 4.2 Approach

Ground truth evaluation for this thesis' in-situ study is approximated by a proxy metric. This proxy metric consists of a number of methods used to evaluate the accuracy of the system. Each method is designed to isolate a certain situation and evaluate its accuracy. These situations included the identification of false positives, identification of true positives, identification of false negatives, the identification of hover events, and the detection of direction of travel.

### 4.2.1 False Positive Identification

Many false positives can be detected by the absence of RFID information. Whenever a crossing is detected by the doorway systems, the corresponding data from the RFID system can be examined to validate the event. If no tag is detected by the RFID system when an event was detected, then that event was not caused by a resident crossing the doorway. False positives detected by the doorway system due to doors, shelves, hinges, and signals reflected from other doorways can therefore be identified as false positives.

### 4.2.2 True Positive Identification

All true positives are identified by checking the corresponding RFID information when a doorway event was detected. Whenever a doorway crossing is detected, the corresponding RFID data can validate a resident's presence near the doorway. Additionally, the RFID system provides individual identification of the person responsible for any event. Since each resident wears bracelets on their ankles with unique tags, they can be identified individually whenever they are near an RFID antenna. This provides a way to evaluate the accuracy of the height measurements made by the ultrasound sensors by matching each height with the individual that was near the door when the height was measured.

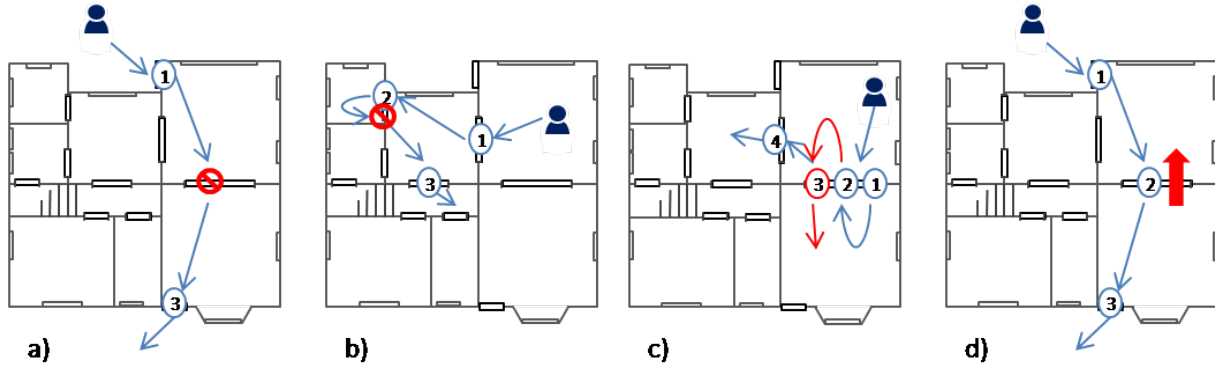


Figure 4.2: False negatives (a,b), hover events (c), and direction (d) are identified by examining the path of residents.

### 4.2.3 False Negative Identification

False negatives were evaluated by leveraging the layout of the home and the knowledge that people cannot teleport from one room to another. First, all doorway crossings not designated as false positives were assigned to the resident identified to be near the doorway when the event occurred by the RFID system. Second, each individual was tracked as they interacted with each doorway. Whenever a person was found to interact with two doorways on opposite sides of a third door without interacting with that third door, a teleport was identified. One example of this situation is shown in Figure 4.2(a). If a scan of the RFID reading at the third door showed the person as near that doorway between the times of the two other events, then a false negative was identified. This method has the limitation of only being able to evaluate transitions where three different adjacent doorways, and therefore 4 different rooms, were contacted.

Events that occurred at leaf rooms, where only one doorway leads into the room itself, and events where the person entered and exited a room through the same doorway, are evaluated by processing the corresponding RFID information. When the doorway system detects two consecutive events that occur at the same doorway the RFID information for that doorway over that period is clustered. If two clusters of RFID readings are detected, indicating that the person walked through twice, then a false negative is recorded as shown in Figure 4.2(b).

### 4.2.4 Resident Behavior Identification

Ideally, doorways are used as a simple passageway, where a person moves from one room to another. This ideal case results in a true *transition*, an event with an identifiable direction and change in room location. However, residents in a real home environment often don't use doorways only to transition from room to room. A person may reach through a doorway to flip off the light in the opposite room or grab and pull closed

an open door. People also often walk into a doorway and stand there, talking to a person in the opposite room without really transitioning into that room. The person may then either enter that room, or return to the room they came from.

When a resident enters a doorway, but never fully transitions into the next room a doorway crossing event as defined by  $D = (time_i, height_t, direction_t)$  doesn't actually apply. These *hover* events inherently have no direction as no transition from room to room occurred. However, the signal processing algorithms attempt to define a direction for every event and will misclassify a hover event as a transition, introducing a false positive into the system. A doorway event must therefore be redefined as either a transition or a hover event, both of which can represent true events that occur in the household.

Hover events are detected and removed from the event set by evaluating all detected crossings where the person should have been detected returning through the door, but was not. This is shown in Figure 4.2(c), where doorway 1, 2, and 3 are contacted in order. To follow the path constraints of the floor plan the person should have returned through the doorway after 2 before continuing on to 3. If the doorway was crossed twice, then clustering all the RFID readings at doorway B between the times when events 2 and 4 were detected should produce 2 RFID clusters. If only one such cluster was produced then event 2 was actually a hover event where the person turned back into the room they came from. This, though not an event that transitions a person from room to room, is not a false positive.

### 4.2.5 Direction Identification

Direction can also be evaluated by following the path of a person through doorways. When three different doorways detect a person consecutively, then the middle event  $(time_i, height_i, direction_i)$  can be compared to this direction. If the estimated direction  $direction_i$  of the event matches the direction found through the path constraints, then the direction was inferred correctly. If it does not, as shown in Figure 4.2(d), then the direction for the event was inferred incorrectly. This method has the limitation of only being able to evaluate direction for internal doorways where there is at least one other doorway on either side of the internal doorway. The house deployed for this experiment had four such doorways.

## 4.3 Validation

To validate this proxy metric, a controlled experiment was run on both doorway systems to be used in the in-situ study.



Figure 4.3: Eleven doorways were instrumented for data collection. Of them, two were external doors and two led to other floors.

### 4.3.1 Experimental Setup

A single-family home with three floors, a main floor, second floor, and basement, was instrumented for this experiment. The 11 doorways on the main floor were instrumented with doorway sensors, including two external doorways, the doorway to the basement, and the stairs to the second floor. The second floor and basement were treated as a single room because no doorways on either floor were instrumented. Data was collected wirelessly from each doorway sensor and logged to a laptop located with the home. This data was then sent to a server at the University of Virginia for analysis.

Ground truth was collected by the same hand-held touch-screen system as describe in Section 3.3.1. To partially emulate an in-situ environment, participants were asked to open and closed doors often as they moved throughout the home. The two participants produced approximately 440 doorway crossings spread evenly over the 11 doorways.

### 4.3.2 Results

Doorway crossings from both in-situ systems (Section 3.1 and Section 5.2 are evaluated using both the proxy metric and ground truth. Precision and recall numbers from each doorway using both ground truth systems is presented in Figure 4.4. The precision and recall of the proxy metric for all doors of the Screen Door System are within 4% of the ground truth numbers. However, Doorjamb 2.0 results in Figure 4.4 show large discrepancies in the precision and recall numbers, up to a 20% difference on all doors. Though distinct trends

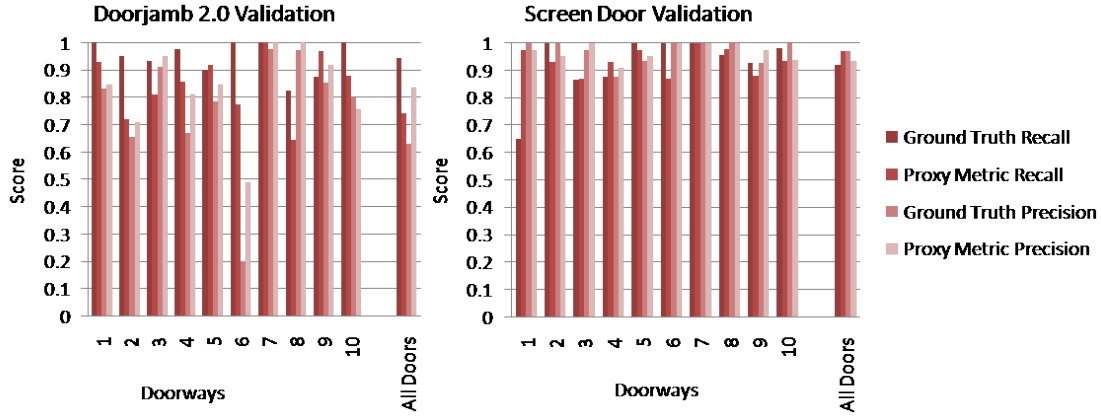


Figure 4.4: The proxy metric better approximated the precision and recall of the Screen Door System, but reflected trends on specific doors of the Doorjamb 2.0 system.

on specific doorways are reflected in the proxy metric, as shown in the low precision of doorway 6, they do not display the same value as the ground truth evaluation.

The higher accuracy of the Screen Door System in general lends itself to produce a more accurate evaluation with the proxy metric. As the proxy metric follows path constraints to identify errors, the low number of false positives in general means that the path is much clearer. This is especially true in the case of leaf room considerations. When the path is clear, it is obvious that a person has walked into and out of a room and RFID does not need to be consulted. If three events are detected when path constraints say either two or four have to occur, then the proxy metric must assign either a false negative or a hover event to the doorway. This assignment is dependent on the crossing detection capabilities of RFID, something that has already been proven to be inaccurate.

Though this issue of hover identification vs. false positive identification can cause errors in the precision and recall evaluation of the proxy metric, it does not invalidate it for use as ground truth. Figure 4.5 shows the F-measure, or harmonic mean, of precision and recall:

$$F = 2 * \frac{precision * recall}{precision + recall} \quad (4.1)$$

For both systems the F-measure obtained by the proxy metric is within 3% of the ground truth evaluation. When examining the F-measure, the decision to assign an event as a false negative or false positive is not greatly felt and the proxy metric provides an evaluation comparable to ground truth.

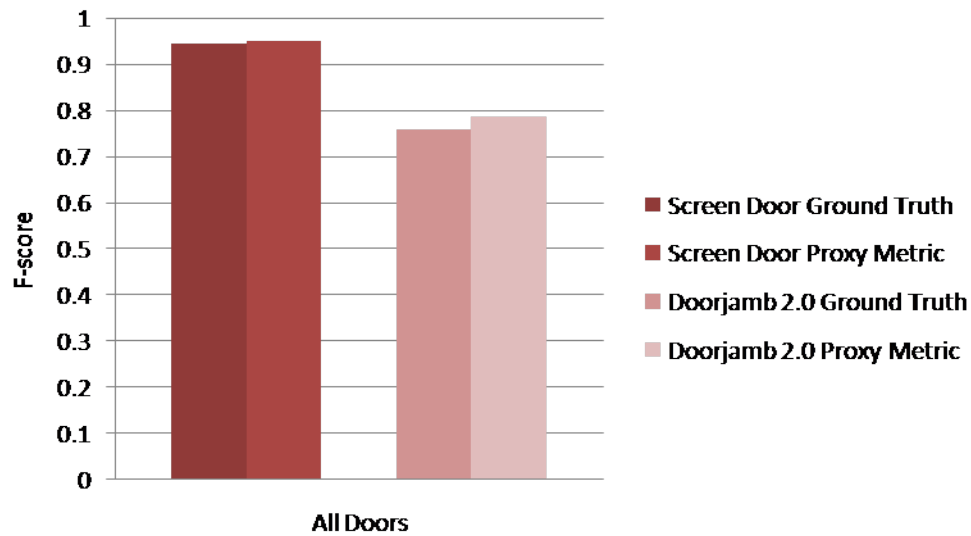


Figure 4.5: The proxy metric differed from the ground truth by less than 3% when the F-measure of all doors was evaluated.

## 4.4 Conclusion

The proxy metric has a number of limitations when used to evaluate a system. First, doorway systems must have generally good recall. Second, the proxy metric relies on two assumptions about RFID: RFID identification can filter out the majority of false positives and can identify two separate true positives when they occur close together at the same door.

The first condition is the major limitation of the proxy metric. This is partially why tracking will not be run on the in-situ data evaluated in this thesis. You will notice from Figure 4.4 that results from only 10 of the 11 doors instrumented with doorway sensors are shown. No data from the 11th doorway sensor was collected during this controlled experiment due to a hardware failure, but the doorway was still crossed approximately 40 times and RFID recorded tag reads at that location. Because the proxy metric relies on path constraints and this 11th doorway was connected to a leaf room, no false negative were detected at that doorway. One of the goals of tracking is to remove false positives while maintaining path consistency for the residents. Because of this tracking can easily remove doorway events at leaf rooms that are true positives and the proxy metric would not detect them as false negatives. Future work will look at creating a ground truth method that can evaluate tracking results.

The second assumption was shown to hold up well for the Screen Door System. However, the high environmental interference experienced by Doorjamb 2.0 data occasionally caused the proxy metric to mis-identify false positives as hover events. Additionally, true positives that occur when a person turns

around and returns through a doorway were occasionally identified as hover events due to insufficient support from RFID information, removing a false negative from the evaluation. Despite these mis-assignments, the F-measure calculated by the proxy metric approximates the ground truth evaluations and general trends of high false positive rates are reflected. For these reasons, the proxy metric will be used to evaluate data from both systems when collected in-situ.



## Chapter 5

# Screen Door System

Doorways and the areas around doorways in a home are used frequently as residents go about their daily lives. Controlled studies will never be able to emulate the variety and frequency of this usage. In order for any tracking system to be used in smarthome applications, the system has to be evaluated in-situ and in the presence of real people.

Doorjamb 2.0, describe in Section 3, has been shown to track residents well in a controlled environment. However, due to some of the physical properties of ultrasonic sensors, particularly the conical shape and range of the sensing region, the accuracy of the system in a real home environment is reduced. This chapter describes a doorway sensor system designed to be used in the environment of a real home and compares these results against Doorjamb 2.0. This system fuses ultrasonic sensors with infrared sensors within a doorway to detect doorway crossings and is called the Screen Door System.

### 5.1 Infrared Range Finders

Infrared range finders measure the distance from the closest object by way of an infrared beam. These devices consist of a transmitter and receiver pair. The angle of reflection calculated when a beam returns to the receiver can be converted into a measure of distance to the reflecting object. A very narrow beam is used to prevent stray light from interfering with the measurements. Because the angle of the reflected light is used to measure distance, objects that are very far away are measured with decreasing accuracy. Additionally, objects that are too close will appear to be far away or not measured. The infrared sensors deployed here measure in the  $15 - 100cm$  range and have a precision of  $1 - 2cm$ .

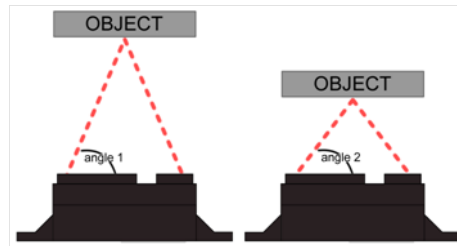


Figure 5.1: Infrared range finders detected the distance to nearby objects by measuring the angle of the reflected light.

Though infrared range finders can measure distance, the precision falls drastically as objects move farther away. The accuracy of infrared reflection also depends greatly on the type of surface being measured. Surfaces such as human hair tend to absorb infrared light, making obtaining height values inaccurate especially in taller doorways or with shorter people. Additionally, due to the narrow beam employed by the sensors, a large number of the sensors would need to be deployed in a doorway in order to ensure the top of a person's head was measured rather than a shoulder. Infrared range finders are also incapable of determining direction due to the narrowness of the beam. For these reasons, infrared range finders alone cannot be used to detect people as they cross doorways.

The narrowness of the infrared beam can greatly increase the precision and recall of detected crossings. The beam is narrow enough that it can be placed in a doorway, pointing straight down, and not detect the surrounding environment. Ideally, it can be placed so that the beam does not come in contact with the door attached to the doorway. Due to slanted doors or imperfect doorjamb this is not always the case, but even then the infrared sensors only detect the door near the floor as a far away object.



Figure 5.2: A screen of infrared sensors ensures that a person crossing the doorway is detected.

To ensure a person is detected when crossing a doorway multiple infrared sensors must be placed evenly along the top of the doorjamb. These sensors form an infrared curtain, or screen, that detects anyone as they

pass through. As the average adult has a head diameter of 6.65 inches or larger, these sensors can be placed so that a person walking sideways is still detected by the screen. A typically sized doorway requires three infrared beams to ensure detection.

## 5.2 Approach

Each doorway sensor within the home contains  $k$  ultrasonic range finders and  $n$  infrared sensors depending on its size. Data collected from these sensors is processed to produce a set of events  $E$  where  $e_i \in E = (door_i, time_i, height_i, direction_i)$ : a 4-tuple with information about what door the crossing occurred at, what time the crossing occurred at, and the estimated height and direction of the person traveling through the door. This process differs from the Doorjamb 2.0 signal processing described in Section 3.1 in two major ways: crossing detection operates only on the infrared data stream and direction is estimated from the fusion of the infrared and ultrasound data streams.

### 5.2.1 Crossing Detection

When a person walks under a doorway one or more of the  $n$  infrared range finders in the doorway sensor will measure a distance. Each infrared sensor is placed 8.5 inches or less from its nearest neighbor forming a screen of infrared beams across the threshold. Though a person crossing the doorway completely perpendicular to the threshold may be able to pass through the screen undetected, this walking pattern is uncommon. Typically, a course grained measurement of the person's head, shoulders, or both will be detected as they cross.

Because the infrared sensors produce a very narrow beam that points directly towards the floor, ideally objects such as doors are not detected even when closed and flush against the doorway. However, doors in residential homes are not always perfectly perpendicular to the floor nor are the doorjambs always perfectly level. Slight imperfections in these areas can cause the infrared sensors to pick up a door when closed, though it does not range far enough outside of the doorway to pick up other environmental objects. To prevent these readings from producing false detections, only infrared readings from the  $n$  infrared sensors whose measured value is above a chosen threshold are considered for event detection. This threshold was set at 4 feet 5 inches above the floor—below the head and shoulders of the average adult, but above the interference of doors. This threshold was also above the height of the resident children in the household and therefore their movements are not captured by the doorway sensing system.

Doorway crossings are detected by scanning all data from the infrared sensors for consecutive measurement events that indicate a head or shoulder. More formally, the  $n$  infrared sensors produce a stream of vectors  $Y$

where each element  $Y^t$  is the readings across all  $n$  infrared sensors at time  $t$ :  $Y^t = y_1^t, y_2^t, \dots, y_n^t$ . Whenever a temporally consecutive vector  $Y^t$  contains a reading,  $y_i^t$ , that passes a threshold all readings from  $Y^t$  are clustered with the previous readings to form a cluster  $C$ . Each cluster  $c_i \in C$  has two vectors  $T$  and  $R$  and a mapping between  $T$  and  $R$  where  $(t_i \in T, r_i \in R)$  is a corresponding infrared reading and timestamp. If a cluster  $C$  has three or more readings above the threshold, the cluster is labeled a detection event. The timestamp of the maximum infrared reading,  $\max(R)$ , is then chosen as the time of the detection event,  $time_i$ .

### 5.2.2 Softening Event Segmentation

When a person walks through a doorway in an in-situ environment they don't always cross and continue on without stopping. The opening and closing of doors becomes a major factor in how a person moves through a doorway. Doors affect the accuracy of the crossing detection when the person and the door they are interacting with are on the opposite sides of the infrared screen. When the person goes to open a door their hand or arm may get picked up by the infrared, then a lower part of their body such as a leg that does not pass the threshold, and finally their head or shoulders as they walk through. When a person walks through a doorway pulling the door closed behind them, or reaching back through to grab and close the door, the same problem occurs.

An aggressively segmenting algorithm would classify these door interaction instances as two separate crossing events. One for the crossing and one for the interaction with the door. However, this would then produce a crossing, where the person returned through the doorway, that didn't occur. To prevent this situation from producing two separate events, the event segmentation is softened. Whenever a reading  $y_i^t$  passes the original event threshold, the surrounding readings in  $y^{t+3}, \dots, y^{t-3}$  are added to the cluster  $C$  if they pass a lower threshold of two feet from the floor. If two clusters both contain these readings they are merged. In this way neighboring events that would usually form two clusters, one a false positive, are merged into one.

### 5.2.3 Height Detection

When a person is directly under the doorway the height measurements are the most accurate, as there is no geometric error produced by their relationship to the ultrasonic sensor. As the infrared only detects an individual when they are directly under the doorway, the infrared cluster  $C$  for an event provides the exact times the height measurements will be most accurate. However, as most of the person's contact with the ultrasonic sensors occurs outside of that narrow infrared window, ultrasound readings  $200msec$

before and after the infrared readings are considered for estimating the height. More formally, the window  $w = [\min(T) - 200\text{msec}, \max(T) + 200\text{msec}]$  is chosen as the time period containing the height of the person.

In addition to the height of the person, the ultrasonic range finders may also pick up the door as a person interacts with it. These measurements often appear taller than the responsible individual. As the location of these doors remains the same as long as a person is not interacting with them, these taller height values remain stable either long before or long after the actual event occurs. Keeping this in mind, a second window  $1\text{sec}$  before and after the window  $w$  is chosen:  $\hat{w} = [\min(w) - 1\text{sec}, \max(w) + 1\text{sec}]$ . If a door has been closed to either side of this event the majority of readings from it should be contained in  $\hat{w} - w$ . Therefore only the height cluster with the lowest ratio of values outside  $w$  is chosen as the height cluster and the maximum height within this cluster is chosen as the person's height. If two clusters have the same noise ratio, the smaller of the two is chosen. If no such cluster exists due to a person walking very quickly, then the measurement closest in time to the event time,  $time_i$ , is chosen.

#### 5.2.4 Direction Detection

Direction estimation relies on the relationship of the physical locations of both the ultrasound cones and the infrared screen. Due to the tilt of the ultrasound sensors, the majority of height readings for any event should be obtained when the person is within range of the cone, on the specific side of the doorway that the ultrasound sensors face. The infrared sensors will pick up the person when they are directly within the doorway. Therefore, depending on the order in which the ultrasound and infrared pick up the individual the direction of travel can be determined. If the majority of height measurements in the event's height cluster occur before the event time, then the direction is estimated to be an entrance, otherwise an exit. If no height was identified for an event then no direction can be determined and is designated as unknown. This direction estimation is combined with  $time_i$  and  $height_i$  to form the final representation of a doorway event  $e_i = (door_i, time_i, height_i, direction_i)$ .

### 5.3 Evaluation

Both Doorjamb 2.0, described in Section 3.1, and the Screen Door System described here were evaluated in an in-situ study. The proxy metric described in Section 4 is used to perform the evaluations.

#### 5.3.1 Experimental Setup

Data was collected over a period of 2 weeks in early March 2013 at the experimental home described in Section 4.3.1. As there was no way to unobtrusively gain information about guests (they would have needed

to provide their actual height and wear RFID bracelets for the duration of their stay) any periods of time identified as containing guests were removed. Though this left 10 days of good data from the doorway system, malfunctions in the RFID bracelets reduced this number. Of the data collected in-situ, 2 days are evaluated here. This data consists of an estimated 3000 doorway crossing.

The two adult occupants participated in the experiment. Their heights were 170cm and 158cm. Each participant wore RFID bracelets around both ankles during the experiment. The home also contains two children, but they were not willing to wear the RFID bracelets for any extended period of time. As the children were too short to be detected by the doorway systems they caused minimal interference.

Individual parameters for each signal processing algorithm was determined using a set of initial controlled experiments attempting to partially emulate an in-situ environment with the use of doors. Once the window parameters and noise thresholds were set they were not adjusted for the in-situ experiment.

### 5.3.2 Results

Each sub component of the event detection from both Doorjamb 2.0 and the Screen Door System were evaluated for accuracy and compared. The conclusion that one system is better than another depends on each component either showing improved results or maintaining the accuracy of the other system.

#### Event Accuracy

Each doorway within a home has a specific environment that affects the precision and recall of events detected there. This environment can greatly affect the precision of doorway sensors as shown in the Doorjamb 2.0 results in Figure 5.3. Doorways with doors attached ( bathroom, bedroom, mudroom) exhibit some of the worst precision. Most of the internal rooms (bedroom, bathroom, hallway, mudroom entrance, upstairs, and nursery) are packed closely together. In these cases precision is lost when ultrasound signals from these surrounded doorways are detected as events.

The Screen Door System is able to filter these environmental issues by using the infrared screen for event detection. Only the frontdoor, surrounded on two sides by doors, has a large number of false positives. The front door is also the only doorway in this study with an arched doorjamb, preventing a perfectly level placement of the sensor. Though the doorway sensor made for the frontdoor has a bottom parallel to the floor where the sensors are placed, this arched frame made attaching the doorway sensor difficult and the final attachment was slightly unstable. This is the likely cause of most of the precision errors on this doorway. Most other doorways in the Screen Door System showed precision around 90% or higher.

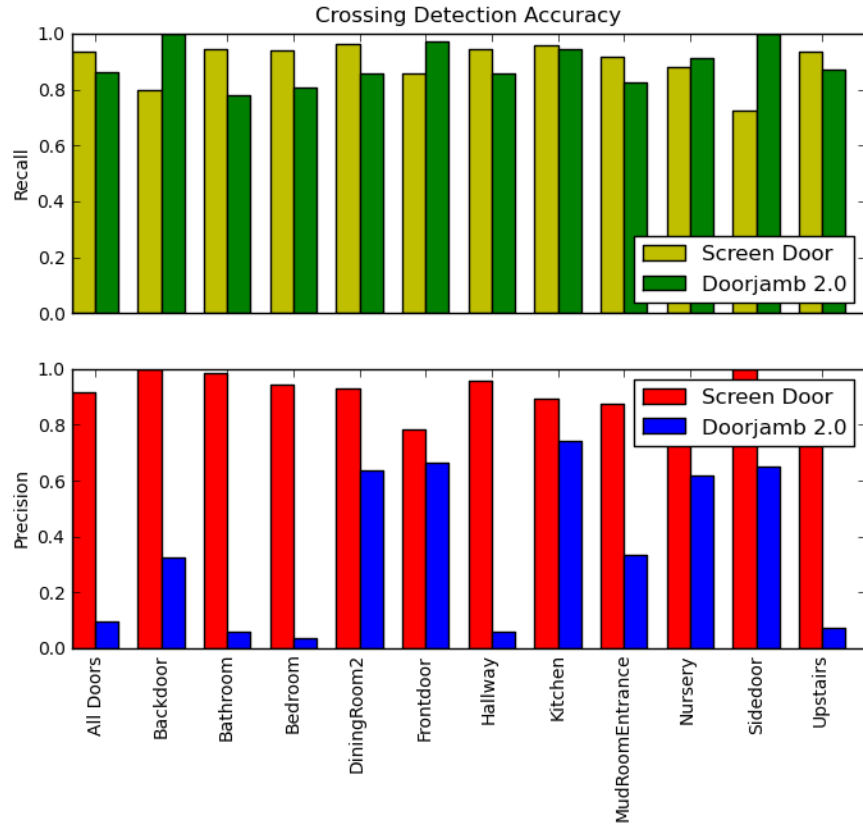


Figure 5.3: Overall the Screen Door System performs better at event detection. The precision of Doorjamb 2.0 in particular is much lower due to the surrounding environmental noise.

Overall, Doorjamb 2.0 exhibited an F-score of 47%. The Screen Door System exhibited an F-score 93%, showing that in the component of crossing detection, the Screen Door System is much better suited for an in-situ environment.

### Height Accuracy

Height is a weak biometric for human identification. Changes in posture while walking, the length of a person's stride, shoes being worn, and the style of hat or hair can all subtly change a person's height. Typically, these changes shorten a person's height. Therefore a person will most commonly appear a few centimeters shorter than their measured standing height, especially in homes where shoes are often not worn. Additionally, even if measurements of height were 100% accurate, these subtle changes would mean that any one person's height should present as a distribution when walking. These changes are also much more common in-situ as residents go about their daily lives.

The accuracy of the height extraction algorithms is evaluated by examining all events where a human

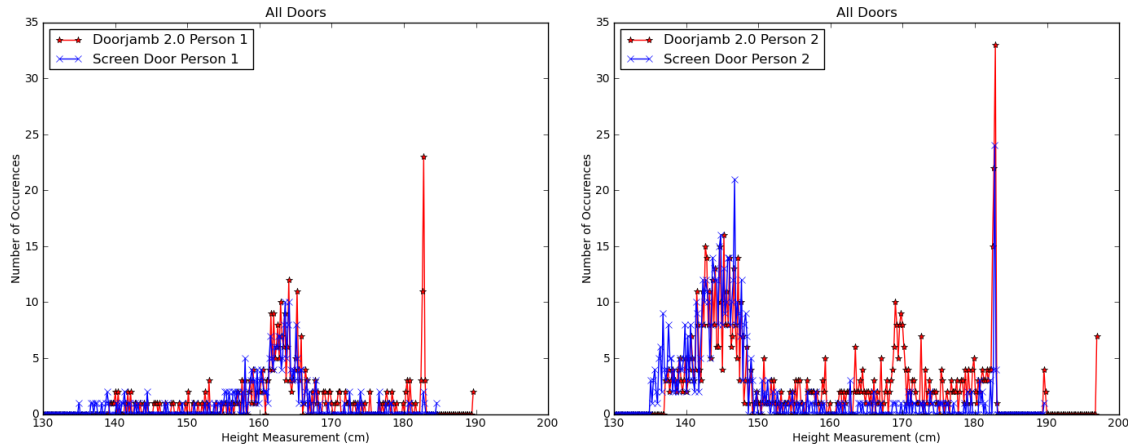


Figure 5.4: Overall, the Screen Door System provides better height measurements for both residents than Doorjamb 2.0.

height was detected. Each event that is matched with ground truth information on the identity of the person responsible is evaluated. Figures 5.4 shows histograms of the height of each person as seen by both of the doorway sensing systems. Overall, both systems produce similar histograms when looking at the heights detected over all doors. However, for person 1 Doorjamb 2.0 identifies some heights much taller than the person as their height. These heights are likely door heights that Doorjamb 2.0 was unable to filter. Person 2 observed this same noise in Doorjamb 2.0 to a lesser degree, but also showed increased noise for a range of 30cm above their actual height. As the height of person 1 is within this range, this means that individual identification will be more ambiguous with Doorjamb 2.0.

Due to the layout of the home, some doors may exhibit more errors in height detection than others. The most damaging case in this experiment was the upstairs doorway. The doorway sensor itself was placed above the stairs, with the ultrasonic cone facing the second floor to avoid interfering with doorway sensors of the doorways immediately in front of the stairs. This placement often caused a person's torso to be measured as they walked up or down the stairs. Due to the geometry of this measurement, it often appeared taller than the person's actual height. If a head height was not measured, or appeared as an outlier, then this higher value was chosen as the estimated height. Doorways with doors also cause some height detection errors for the same reason. Many such errors are filtered by the detection algorithm, but in cases where when a door is open and closed quickly within an event window there is no extensive surrounding data of a similar height to help filter it out for either system. Both systems showed the effect of this environmental noise.

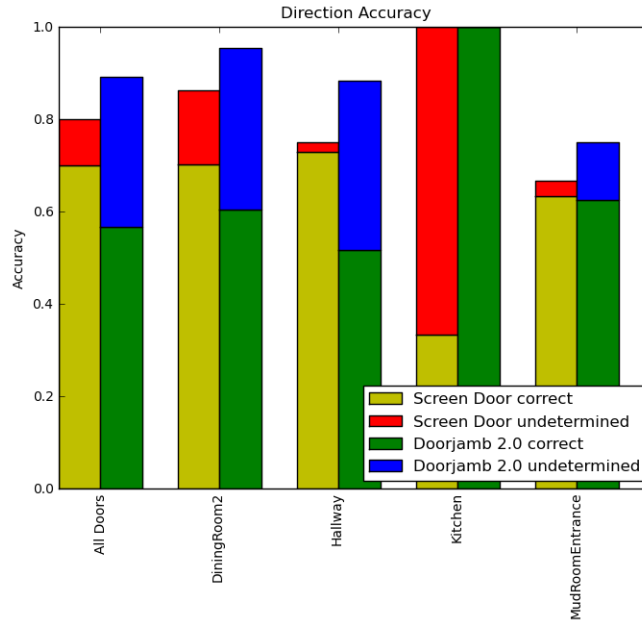


Figure 5.5: The Screen Door System can determine a direction more often and estimate it more accurately than Doorjamb 2.0.

### Direction Accuracy

Direction determination for both systems relies heavily on the accuracy of the event and height components. If an event is detected with no height or only one height value, direction can be either impossible to determine or likely inaccurate.

Given the ground truth challenges encountered in data collection, the direction accuracy of the system is only evaluated on the 4 internal doorways on the home. The results presented in Figure 5.5 show the percentage of correctly identified directions when a direction was identified and the percentage of times the systems were unable to determine a direction. Doorjamb 2.0 was unable to determine the direction of a crossing much more often than the Screen Door System and showed lower accuracy when the direction could be estimated. This is because Doorjamb 2.0 must determine the direction from ultrasound measurements alone, instead of relying on the infrared data to supply the time a person is exactly within a doorway. Though the Screen Door System was able to detect direction more often than Doorjamb 2.0, it also got a higher percentage of the detected directions incorrect. This is because if the time of an event is even slightly incorrect, the Screen Door System may flip its direction estimate. More accurate event times found by the crossing detection component would improve direction detection for the Screen Door System.

## 5.4 Conclusion

In a real home environment the Doorjamb 2.0 tracking system is unable to accurately detect people as they walk through doorways. Though an in-situ study, the Screen Door System that fuses data from these two types of sensors is shown to accurately detect doorway crossings and more accurately estimate the height of the responsible person and their direction of travel. The Screen Door System address the in-situ challenges of doors, surrounding furniture, interference from other ultrasound sources, and residents walking by but not through doorways. This accuracy will allow future smarthome applications that leverage this tracking information to operate in a real home environment.

## Chapter 6

# Smarthome Applications

Tracking residents is the foundation for many future smarthome applications. This chapter discusses a case study of one such application. In the future, tracking will be used to improve a variety of existing applications, such as HVAC and light control, as well as provide a foundation for new smarthome technologies and ideas.

### 6.1 Case Study: Floor Plan Inference

The installation of many smarthome sensing systems requires information about the floor plan of the building. In some cases, floor plans are impossible to obtain for residential homes and homeowners often want to avoid time intensive installation procedures that require the information to be collected manually. Prior work has shown the floor plan of a home can be generated using an array of motion sensors placed throughout a home [23]. These motion sensors can then be clustered into a series of adjacent rooms based on when and which sensors detect residents moving around. Because these motion sensors can see between rooms through an open doorway, the end result of this process is a small set of floor plans for a homeowner to choose from.

Using the doorway sensing system, the floor plan of a home can be generated after a sufficient amount of collective doorway crossings  $E$  have been observed. This floor plan generation algorithm uses a similar concept to [23], but requires less aggregate data to produce an accurate floor plan because the doorways provide the boundaries for each possible room. This system can also result in a single accurate floor plan rather than a set of candidates. This work uses the Screen Door System and nine days of collected in-situ data to generate the floor plan for our deployed household.

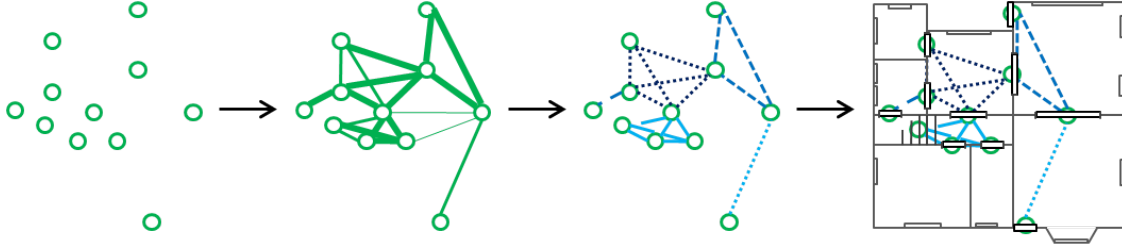


Figure 6.1: Rooms are formed by doorway sensors with no initial location information through two steps. First, an adjacency graph is formed based on the movements of residents within the home. Second, these adjacency edges are filtered to remove noise and clustered to form completely adjacent cliques of doorway sensors bordering rooms.

### 6.1.1 Approach

Two steps are used to infer the floor plan of a home. First, the adjacency matrix between rooms is generated from the doorway data. Second, a floor plan is generated based on the window and doorway direction constraints by the same process used in [23]. The first step is the only step modified by this work and is described below.

#### Room Adjacency Generation

The layout of doors within a home and the movements of residents follow a few concrete rules that can be leveraged to generate a floor plan. These rules do not change across residents or households.

1. Residents can only travel between rooms through a doorway.
2. Doorways bordering the same room are walked through in succession, when a resident enters and exits.
3. A door must border two rooms.

Applying rule 1, we know that each new event  $e_{i+1}$  is from a doorway adjacent to that of the previous event  $e_i$  because the resident walked through no other door in between the two. Each event  $e_i$  can therefore be paired with its two temporal neighbors to form a set of links between doors  $(d_i, d_{i+1}), (d_{i+1}, d_{i+2}), \dots \in L$ . This data can then be represented by a graph  $G$ , as shown in Figure 6.1, where there are  $k$  nodes representing the  $k$  doorway sensors and each link  $(d_i, d_{i+1})$  adds a weight of 1 to the corresponding edge. Edges that loop back to the original nodes, when someone enters and exits a room through the same door without traveling to another room in the between, are not included.

Internal rooms can be generated from  $G$  by applying rule 2. Any doors that are walked through in succession border the same room, and any two nodes in  $G$  that are connected by an edge have been walked through in succession. Therefore, any two nodes in  $G$  that share an edge also share a room. This information

can be leveraged to create rooms from the graph of adjacent doorways  $G$ . For example, assume room  $r_1$  has three doors  $d_1, d_2$  and  $d_3$ . Each of these doors are adjacent to each other by virtue of bordering the same room. Therefore, each doorway in  $G$  will be connected by an edge as residents walk from one door to another through  $r_1$ . Hence, the subgraph made from  $d_1, d_2$  and  $d_3$  is fully connected and, when examined from the perspective of the whole graph  $G$ , is a clique. To prevent a subset of the doors  $d_1, d_2$  and  $d_3$  claiming to border a room other than  $r_1$  the maximal clique must be chosen.

Each clique in  $G$  represents an internal room in the floor plan that has two or more doors. Leaf rooms, or rooms with only one door, can be generated from these cliques by applying rule 3. Any door, or node in  $G$ , that is not connect to two internal rooms is also connected to a leaf room. Hence single-door rooms, such as the majority of bathrooms, can be generated from  $G$ .

### Noise Filtering

Up to this point we have assumed that rule 1, residents can only travel between rooms though a doorway, is always followed. Physically this is still true, residents cannot teleport. However, we cannot assume that the doorway sensors are 100% accurate. False negatives, when a resident travels through a door but is not detected, and false positives, where an event is detected but no resident has changed rooms, can produce links between doors in  $L$  that would not exist in the actual floor plan. These false links occur infrequently compared to the true links generated by doors that are actually adjacent. Therefore, removing any links below a certain threshold, a weight of 18 for our deployment, removes all false edges from graph  $G$ .

However, we must realize an assumption that's been made. We've assumed there is only one resident traveling throughout the household. When two residents are within the home, as was true for much of our deployment, two successive events may be made by two different residents in two non-adjacent doorways. Therefore some links in  $L$  formed when two residents are present do not suffer from false negatives or positives, but also do not represent true doorway adjacency. We found this was especially frequent in the doors of the two most commonly used rooms in our deployment, the bathroom and the kitchen. Though those two rooms were separated by three other rooms in the household their doors exhibited stronger edges than many true adjacent doors.

To filter out false links caused by two residents we leveraged the identity information of height from each event  $e_i$ . Links in  $L$  where  $h_i - h_{i+1} < 10cm$  were removed. 10 cm was chosen as it was the height difference of our participating residents. Some true adjacencies were lost due to this filter, such as when a resident stooped when going through a subsequent doorway or otherwise changed their height, but this loss was relatively small. In general, the links removed from  $L$  were those generated by two different residents. The

graph  $G$ , once also filtered by the false negative and positive event threshold, holds only edges that represent true doorway adjacencies.

### 6.1.2 Evaluation

Data collected from the same deployment as Section 5.3.1 was used to evaluate this smarthome application. As the ground truth of this algorithm is the floor plan of the house itself, ground truth information about the location of residents was not needed. Therefore, data collected without corresponding RFID information was used to evaluate this application. In total, nine days of data collected over the two week period were used.

The room adjacency graph  $G$  was produced for the deployed household with 100% accuracy. All rooms were generated and placed adjacent to the matching rooms in the physical floor plan. Each external room, from both the back door of the household and the front door were also correctly allocated. However, the algorithm is unable to identify that these are "outside" areas. Some information outside of the data collected by the doorway sensors is needed to make this distinction. This is easily done by marking these 2 out of 11 sensors as monitoring external doors. Discovering if this marking is possible without user input is an avenue of future work.

The adjacency graph  $G$  can then be run through the search algorithm for floor plans as described in [23]. Since this new room adjacency matrix generation takes less time than that of the prior work [23], on the order of days as opposed to weeks, a floor plan for the home can be generated much more quickly. This means that any home that requires floor plan generation for smarthome applications can be configured much more quickly.

### 6.1.3 Conclusion

This case study does not make the claim that floor plan generation from doorway sensors is possible in all homes. Each home and its residents are different and may present obstacles to accurate floor plan generation that were not seen here and have not been anticipated. However, this case study shows an exciting possible smarthome application that the doorway sensing system can provide the foundation for.

## Chapter 7

# Conclusion

Doorway sensors present a promising solution to tracking residents in a real home environment. They are non-invasive and require little to no participation from residents in order to perform tracking. Many smarthome applications can be built on top of resident location information, including building knowledge of a home's floor plan from doorway sensors alone. However, the noisy daily lives and environments of real people must be taken into consideration when designing a tracking system for use in real homes. Of the two hardware platforms presented here, only one was able to handle the noise of in-situ environment. Though Doorjamb 2.0 performed well in a controlled environment, it was unable to filter the environmental noise found in a real home even when redesigned to do so. On the other hand, the Screen Door hybrid system held up well in the presence of environmental noise. With a system designed to track residents in a home environment like the Screen Door System, the possibility of new and exciting smarthome applications that can operate in homes becomes feasible. HVAC and lighting control can be customized to a specific individual, appliance use in a home can be assigned to specific residents and billed accordingly, and the wasted energy that is not consumed by the residents of the home can be saved.

### 7.1 Limitations and Future Work

As this thesis does not perform tracking on the data collected in-situ, future work will be on the evaluation of the Screen Door System when fed through tracking. Ground truth collection issues for the in-situ experiment currently prevents the use of tracking to improve the results obtained by both doorway systems. As such, room occupancy detection performed by tracking cannot be evaluated, nor can the possible improvements to precision and recall that are a benefit of the tracking algorithm. Therefore, alternate forms of ground truth,

or improvements to the RFID system that will make it accurate enough for use as ground truth, are the main focus of future work.

The presented tracking system can be improved in a number of ways. Firstly, the hardware platform providing data to the detection algorithms could be improved. The doorway sensors themselves contain a number of sensors that could be used to improve tracking, most notably the PIR motion sensors. These sensors could provide the system with further information to about the location of residents within the home, like helping to confirm which side of the door a resident is on. Additionally, improvements to the data collection rate of the ultrasonic sensors could be leveraged to better measure and identify individuals. Secondly, work can be done to allow the sensors to be more easily used in residential home. The number of people in the home and their heights could be learned by the system when the sensors are first installed. Guests could be identified by the system when they enter and leave the home and noise parameters for each door could be learned in-situ.

A residential tracking system also paves the way for new smarthome technologies. HVAC control, fixture use identification, fixture location identification, and more detailed floor plan generation are all possible future applications of the tracking system.

# Bibliography

- [1] American time use survey. 2010.
- [2] T.W. Hnat, E. Griffiths, R. Dawson, and K. Whitehouse. Doorjamb: Unobtrusive Room-level Tracking of People in Homes using Doorway Sensors. In *SenSys '12: the 10th ACM Conference on Embedded Networked Sensor Systems*, pages 309–322. ACM, 2012.
- [3] V. Srinivasan, J. Stankovic, and K. Whitehouse. Using height sensors for biometric identification in multi-resident homes. *Pervasive Computing*, pages 337–354, 2010.
- [4] R. Want and A. Hopper. Active badges and personal interactive computing objects. *Consumer Electronics, IEEE Transactions on*, 38(1):10–20, 1992.
- [5] P. Bahl and V.N. Padmanabhan. Radar: An in-building rf-based user location and tracking system. In *INFOCOM 2000. Nineteenth Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings. IEEE*, volume 2, pages 775–784. Ieee, 2000.
- [6] Konrad Lorincz and Matt Welsh. MoteTrack: A Robust, Decentralized Approach to RF-Based Location Tracking. In *LoCA*, volume 3479 of *Lecture Notes in Computer Science*, pages 63–82. Springer, 2005.
- [7] A. Harter, A. Hopper, P. Steggles, A. Ward, and P. Webster. The anatomy of a context-aware application. *Wireless Networks*, 8(2):187–197, 2002.
- [8] N.B. Priyantha, A. Chakraborty, and H. Balakrishnan. The cricket location-support system. In *Proceedings of the 6th annual international conference on Mobile computing and networking*, pages 32–43. ACM, 2000.
- [9] K. Chawla, G. Robins, and L. Zhang. Efficient rfid-based mobile object localization. In *Wireless and Mobile Computing, Networking and Communications (WiMob), 2010 IEEE 6th International Conference on*, pages 683–690. IEEE, 2010.
- [10] L.M. Ni, Y. Liu, Y.C. Lau, and A.P. Patil. Landmarc: indoor location sensing using active rfid. *Wireless Networks*, 10(6):701–710, 2004.
- [11] M. Buettner, R. Prasad, M. Philipose, and D. Wetherall. Recognizing daily activities with rfid-based sensors. In *The 11th international conference on Ubiquitous computing*, pages 51–60. ACM, 2009.
- [12] A.B.M. Musa and J. Eriksson. Tracking unmodified smartphones using wi-fi monitors. In *Proceedings of the 10th ACM Conference on Embedded Networked Sensor Systems*, pages 281–294. ACM, 2012.
- [13] W. ur Rehman, E. de Lara, and S. Saroiu. Cilos: a cdma indoor localization system. In *Proc. UbiComp*, pages 104–113.
- [14] 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. In *Proceedings of the 9th ACM Conference on Embedded Networked Sensor Systems (SenSys)*, 2011.
- [15] D. Gafurov and E. Sneekenes. Gait recognition using wearable motion recording sensors. *EURASIP Journal on Advances in Signal Processing*, 2009:7, 2009.

- [16] Z. Liu and S. Sarkar. Outdoor recognition at a distance by fusing gait and face. *Image and Vision Computing*, 25(6):817–832, 2007.
- [17] T. Teixeira and A. Savvides. Lightweight people counting and localizing in indoor spaces using camera sensor nodes. In *Distributed Smart Cameras, 2007. ICDSC'07. First ACM/IEEE International Conference on*, pages 36–43. IEEE, 2007.
- [18] M.D. Addlesee, A. Jones, F. Livesey, and F. Samaria. The orl active floor [sensor system]. *Personal Communications, IEEE*, 4(5):35–41, 1997.
- [19] Y.L. Shen and C.S. Shin. Distributed sensing floor for an intelligent environment. *Sensors Journal, IEEE*, 9(12):1673–1678, 2009.
- [20] J. Schiff and K. Goldberg. Automated intruder tracking using particle filtering and a network of binary motion sensors. In *Automation Science and Engineering, 2006. CASE'06. IEEE International Conference on*, pages 580–587. IEEE, 2006.
- [21] N. Patwari and J. Wilson. Rf sensor networks for device-free localization: Measurements, models, and algorithms. *Proceedings of the IEEE*, 98(11):1961–1973, 2010.
- [22] Yao Y. Ranjan, J. and K. Whitehouse. Rf doormat: Room-level tracking in homes using rfid. *under submission*, 2013.
- [23] Jiakang Lu and Kamin Whitehouse. Smart blueprints: Automatically generated maps of homes and the devices within them. *The 14th ACM International Conference on Ubiquitous Computing (Ubicomp'12)*, 2012.