

Using Height Sensors for Biometric Identification in Multi-resident Homes

Vijay Srinivasan, John Stankovic, and Kamin Whitehouse

Department of Computer science, University of Virginia,
Charlottesville, Virginia, USA
`{vs8h,stankovic,whitehouse}@cs.virginia.edu`

Abstract. In this study, we evaluate the use of height for biometric identification of residents, by mounting ultrasonic distance sensors above the doorways in a home. Height sensors are cheap, are convenient for the residents, are simple to install in an existing home, and are perceived to be less invasive than cameras or microphones. Height is typically only a weak biometric, but we show that it is well suited for identifying among a *few* residents *in the home*, and can potentially be improved by using the history of height measurements at multiple doorways in a tracking approach. We evaluate this approach using 20 people in a controlled laboratory environment and by installing in 3 natural, home environments. We combine these results with public anthropometric data sets that contain the heights of residents in 2077 elderly multi-resident homes to conclude that height sensors could *potentially* achieve at least 95% identification accuracy in 95% of elderly homes in the US.

1 Introduction

The ability to identify residents in a home is crucial for many smart home applications: in order to respond to activities in the home, the system must be able to identify *who* is in a particular location or performing a particular action such as cooking or exercising. Our recent discussions with a commercial peace of mind elderly monitoring enterprise [4] reveal several interesting *user requirements* for accurate, long term resident identification and tracking in homes: (1) residents will not wear tags or manually identify themselves at every room for long periods of time, (2) residents will not allow 'perceived' invasive devices such as cameras or microphones in the home, and (3) residents want the sensors to be fairly invisible, similar to existing motion sensor installations, and do not want an expensive building retrofit.

Existing *implementations* that perform resident identification have several advantages, but fall short of our set of requirements above. Some approaches are inconvenient because they require the user to *wear a tag* [13, 21], or to *manually* trigger a biometric sensor such as a thumbprint or retina scanner [18]. Some systems require cameras for gait, form, or face recognition [16], but cameras are often *perceived as invasive* because they can be used to collect much more information than just the user's identity [11]. Other implementations require

structural changes to the home, such as instrumenting the floor [6,9] with force plates, which can incur high cost and effort. Many practical smart home applications such as in-home medical care for the elderly [17] and occupant-based energy monitoring [8] cannot use solutions that inconvenience the user, are intrusive, or require an expensive building retrofit. For this reason, some commercial deployments of smart home technologies today are limited to single-resident homes or do not fully monitor information about multiple residents [2,4].

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

The main contribution of this work is to demonstrate that height can be effective for biometric identification *in the home*. We evaluate the use of height as a biometric in four ways: (1) We quantify the biometric error of our approach using 20 subjects in a controlled laboratory environment, in which we vary the direction, speed, and location of the person walking under the doorway. (2) We measure the degree to which height sensors can identify room occupancy of residents in 3 natural home environments for 5 days each. (3) We use public anthropometric data containing the heights of elderly residents in 2077 multi-resident homes from the 2006 health and retirement study to estimate that our approach can *potentially* achieve at least 95% identification accuracy in **85%** of elderly US homes sampled in this study. (4) Through a simulation study, we show that incorporating the history of height measurements at multiple doorways using a tracking approach can *potentially* increase the proportion of homes where our solution is applicable with 95% accuracy from 85% to 95%, and also reduce the height difference required for 99% identification accuracy from 7cm to 3.25cm. We quantitatively compare our approach against two other state of the art non-invasive resident identification implementations, namely anonymous binary sensor and activity model based multi-resident tracking [20], and weight sensing [9], and find that our approach achieves improvements in identification

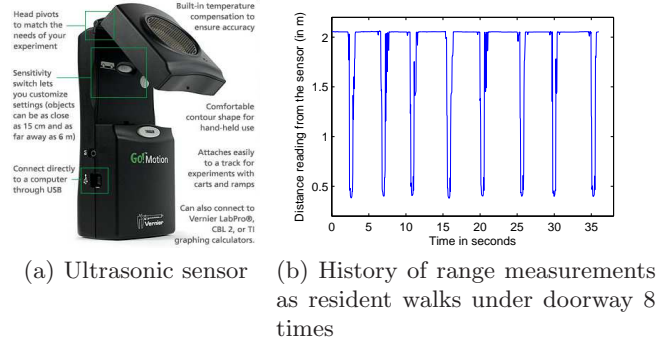


Fig. 1. Our study used the Go Motion ultrasonic range finder mounted above doorways (a). As users walked beneath the sensor, the range measurements changed (b).

accuracy or installation effort and cost compared to *existing implementations* of these approaches.

The rest of the paper is organized as follows. Section 2 discusses existing resident identification solutions from the literature. Section 3 gives an overview of our approach and describes our algorithm to sense height from ultrasonic distance measurements. Sections 4 and 5 describe the results of our controlled lab experiments and natural, in-home experiments respectively. Section 6 analyzes how our empirical results potentially extrapolate to a national level using public anthropometric data. Section 7 evaluates the potential improvement in the inherent identification accuracy of the height sensor by tracking the history of height measurements at multiple locations in a home. Section 8 discusses the application of height sensors for in-home room level tracking, and systematically lists the advantages and limitations of our current approach and study. We conclude by summarizing our findings in section 9.

2 Existing Solutions for Resident Identification

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

Tag-and-track approaches operate by placing a uniquely identifiable device on each individual resident. This approach has been widely used since the Active Badge system almost two decades ago [19], and in other systems since. The pedestrian localization system proposed by Woodman et al [21] uses a foot mounted inertial sensor to track pedestrians to within .7 meters 95% of the time in a large office building with no additional infrastructure. More recently, innovative tracking solutions that require a very low infrastructure cost [12, 15] are emerging. Tag and track approaches have three important advantages: (a) High location granularity with little or no infrastructure, (b) Selective preservation of

location privacy by switching off device, and (c) Highly scalable with respect to the number of residents in an indoor space: each user can be given a new device with a uniquely identifying number. However, one drawback of this approach is that it requires the user to actively carry the device at all times when location information is desired. It can be an inconvenience to in-home residents for long-term deployments. In our past experience with deployments, and while trying to use MoteTrack [13] wireless tags to collect ground truth in this study, users frequently forget to carry their tag, especially immediately after waking up or showering. Automatically reminding residents to carry the device is an option, but we believe such an approach is intrusive and inconveniences the user for long-term applications such as elderly medical monitoring.

Several indoor resident identification systems use **cameras** for computer-vision based face, shape, and gait recognition [16], or **microphones** and audio signal processing. These approaches might require expensive on-board computation or high communication bandwidth to a central base station that executes the vision algorithms, but are passive and highly accurate. However, user studies by researchers from Intel and companies like WellAware have found that a large fraction of potential users have *perceived* privacy concerns about cameras or microphone sensors [11]. Therefore, this class of approaches is most appropriate for short-term situations in which rapid deployment, and/or high accuracy are important, and where long-term privacy concerns of monitoring residents in their own homes are not an issue.

Wilson et al, in 2005, propose using only resident **usage models of anonymous motion sensors in rooms and switch sensors on daily-use objects**, and **resident activity models** to identify and track their activities and locations [20]. They propose using a particle filter that uses Markov state transition and sensor use models learned from short term training data, obtained using a tag and track approach or manual labeling. The main advantage of this approach is that the simple single-pixel sensors are cheap and easy to install, and are not perceived to be invasive or inconvenient. However, an important drawback of this approach is low accuracy: this system was reported to have 70% accuracy when tracking 3 residents over a week-long period, and in our own deployments in 3 multi-resident homes, we observed this approach to have accuracies of 65-75%. These accuracy rates may be reasonable for some applications, but confusing the identities of residents more than a third of the time could cause problems for some smart home applications such as medical monitoring. To increase identification accuracy, additional biometric sensor data, as discussed in this paper, can be included in the STAR particle filter.

Several systems including the Active Floor and the Smart Floor **instrument the floor** to locate and identify individuals [6, 9]. Jenkins et al, in 2007, studied the effectiveness of using resident mass, derived using force plate signals, to identify multiple individuals in a large population [9]. Gait analysis can also be used to differentiate individuals from instrumented floors. This type of single-pixel mass-based identification approach has the advantage that it can be performed without inconveniencing the user or violating resident privacy. Existing force

plates and smart floors require careful installation to improve aesthetic appeal and achieve user acceptance; accurate force plates are also very expensive. *However*, more *compact, cheap designs* of weight sensors that have the same form factor as a floor mat can be explored for easier installation and better aesthetic appeal in the home.

Height is a weak biometric that is often used on driver’s licenses or police reports, and it can not be used to definitively identify individuals from a large population. Some existing systems use invasive video cameras to identify height [5]. Nishida et al [14] propose instrumenting the entire ceiling with a dense set of ultrasonic devices to perform fine-grained location tracking with an ultrasonic radar system. However, this system is not evaluated for its ability to differentiate or identify individuals, and this approach would involve substantial deployment effort. In 2006, Jenkins et al [10] proposes placing infrared or ultrasonic distance detectors on top of doorways for identification based on height (in a poster), using an approach similar to that described in this paper. However, height sensors are not experimentally evaluated for accuracy, multiple readings are not combined as the user walks through the home to improve accuracy, and the poster does not analyze the wider ramifications of height sensing on *in-home* resident identification. To the best of our knowledge, our work is the first to analyze how height sensing can potentially be used to effectively address the multi-resident room location and identification problem in homes with high accuracy.

3 Overview: Sensing Height with Ultrasonic Sensors

To identify residents as they move throughout a home, we deploy an ultrasonic distance sensor above every exit and entry into a room. We used an off-the-shelf ultrasonic distance sensor [3] shown in figure 1a. This distance sensor sends out ultrasonic pulses at 50KHz in a diverging cone 15 to 20° off the axis of the beam. The device then measures the time taken for the echo to return, and uses it to calculate the distance to any obstacle in front of it. Only the minimal distance of any obstacle is reported. For example, when a resident stands under the device, only the distance to the top of the head is reported, while distances to the shoulder, ear etc are automatically filtered out.

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

Our algorithm to extract height events and height values is as follows. We first compute timestamps when the reported distance is below the default distance with no obstacles. We then cluster these timestamps using the DB-SCAN clustering algorithm [7] to compute discrete **height events**, that correspond to residents passing by or standing under the sensor. This clustering process eliminates most noise due to a single, spurious reading. Then, for each cluster of low readings, we find the minimum distance reported (i.e. the maximum height value measured). We subtract that measurement from the default height measurement with no obstacles and use the result to be the height measurement for that height event.

To *identify* residents based on measured height values, we use a Maximum Likelihood Estimate (MLE) classifier to assign each height event to one of multiple candidate residents in a home. For each height event, the MLE classifier simply computes the probability that each resident triggered it, based on the height of that resident and the error distribution of the sensor, and assigns the height event to the resident that maximizes the likelihood of the observed measurements. In the next section, we collect height data from 20 test subjects using controlled experiments in a lab to characterize the error distribution of height measurements, under diverse scenarios of passing through a doorway.

4 Experiments in a Controlled Lab Environment

4.1 Experimental setup

We characterize the error distribution of height sensors in a controlled lab setting by placing the ultrasonic sensor on top of a doorway about 90 cm wide and having 20 users with known heights pass beneath the sensor in a controlled manner. We chose a doorway of this width because it matched the width of many of the doorways seen in our real home deployments. We selected 20 subjects of differing heights for our experiment. The distribution of heights among the 20 subjects can be inferred from any of the scatter plots in figure 2. The subjects were randomly chosen from a pool of graduate students from 20-30 years of age; 16 of our subjects were male and 4 were female.

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

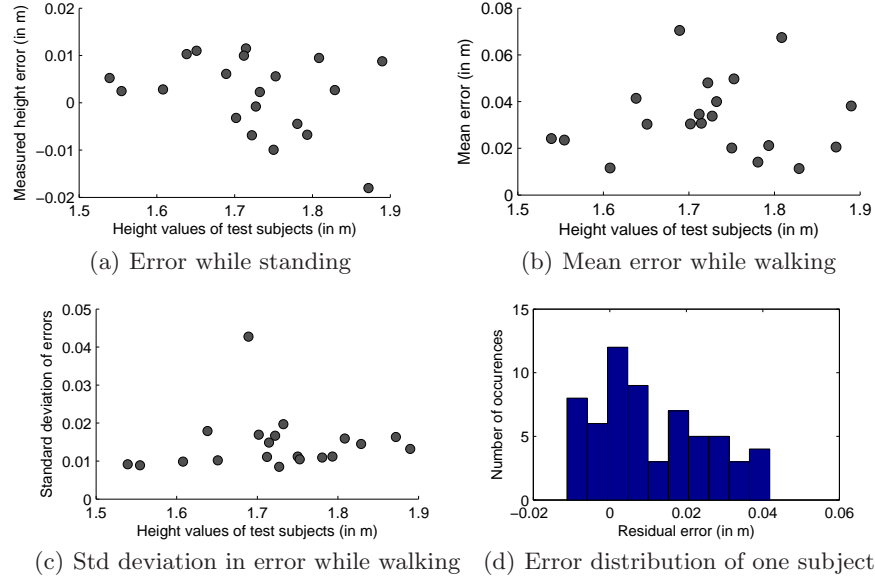


Fig. 2. Controlled laboratory experiments indicate low measured error while standing. Mean error while walking is higher due to a natural reduction in height compared to standing erect, and different walking styles. Standard error while walking is very low. The error distribution approximates a log normal distribution.

4.2 Evaluation results

Figure 2a illustrates that, when residents are standing erect beneath the sensor, the average error across all 20 subjects is only 0.2cm, and the maximum height error is 1.15 cm. Figures 2b and c show scatter plot of mean height measurement error and standard deviation in error while *walking*. *Error* here refers to the difference between the manually measured height and the height output by our height based identification algorithm for each height event. The results shown in figure 2 use the aggregated height data from all our walking experiments without shoes. We do not include results with shoes here, but observe that the mean measured height simply increased by the height of the shoe on average, and changes to the standard deviation of errors were negligible with shoes on.

From figure 2b, walking height as measured by our sensors is lower than erect, standing height by 3.31 cm on average across all subjects. This is possibly due to the natural decrease in *apparent height* as a person walks. Also, different walking styles such as bending and keeping heads down contribute to this decrease compared to erect standing height. More important is the *deviation* in residual error, the standard error, for each subject across different height events, since this will be crucial in determining identification accuracy among multiple residents in a home. We note that the mean deviation in error is only 1.45 cm. This low deviation implies that 99% identification accuracy can be obtained as

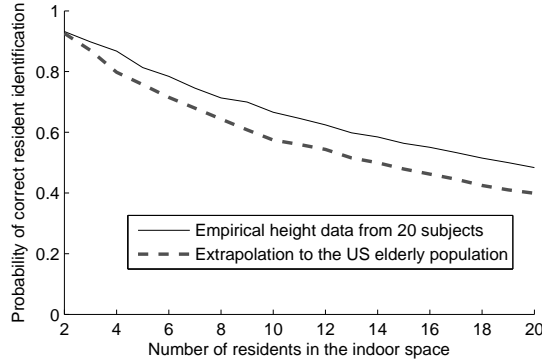


Fig. 3. Height measurements become less effective for biometric identification as the pool of individuals increases. The heights in our study were easier to differentiate than those of the general population.

long as the heights of two residents are 7cm apart. We explore this tradeoff more fully in section 7, when we describe our history based tracking algorithm using height.

In figure 2d, we show an example distribution of residual error from one test subject. The distribution shown here suggests a log normal distribution rather than a normal distribution for residual error. Thus, we ran hypotheses tests at .03 significance level for all subjects to test two different hypotheses (a) distribution of height values is normal (b) distribution of height values is log-normal. *The proportion of subjects for which the normal and log-normal hypotheses could not be rejected are 75% and 85% respectively.* The log-normal distribution, which skews naturally to the right, appears to be a better fit for modeling measured height. This is because the apparent height of a person very rarely *increases* (perhaps due to thick shoes) while walking, but more often decreases in experiments due to the ultrasonic beam making contact with the person’s side (e.g.) shoulder or ear) instead of the head.

Using the empirical height data collected in the lab, we compute how well height can differentiate among a fixed set of N residents in an indoor space. In particular, we empirically calculate the accuracy with which height events are assigned to their ground truth test subjects. In figure 3, we show how this empirical identification accuracy using height decreases as we increase the number of residents under consideration in the indoor space; we randomly choose residents from our pool of 20 subjects, and evaluate how accurately individual height events generated by our subjects are labeled using a log normal MLE classifier trained from the controlled experiment data. For each N value, we repeat the random sampling 100 times. Also shown in figure 3 is how this analysis extrapolates to a national level. We model the mean and standard deviation in residual height error as a function of height using two types of curve fit models: simple linear curve fitting, and nearest neighbor interpolation. Thus, given the

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

The 2006 health and retirement study [1] (HRS 2006) contains height measures of 4154 elderly residents living in multi-resident households. We randomly sample a fixed number N from this set of residents to be identified in an indoor space. We then *analytically* calculate the probability that any height event is assigned to the correct resident among the N residents, assuming a MLE classifier that uses the log normal distributions derived from our nearest neighbor and linear fit models. In figure 3, we show how this probability of correct event labeling degrades as we increase the number of residents in the home, randomly sampling 100 different sets of elderly residents for each N value. We only show the results with the nearest neighbor curve fit model, since only negligible differences were observed when the linear fit model was used.

As we can see in figure 3, for indoor spaces strongly resembling 2 or 3-resident elderly homes, height based identification has a mean accuracy of 87-92% using both empirical data and extrapolation to the national level. In particular, we note that 99% of the elderly multi-resident households with valid height measurements in the national study were 2-resident households. We use this insight in the next section to demonstrate the high accuracy of height based identification in 3 real multi-resident home deployments. For households with 4 residents, the identification accuracy drops to 77%; in section 7, we show that by using the history of height events at multiple doorways in a home, we can improve the identification accuracy in even 4-resident homes to 90%.



Fig. 4. Height sensors deployed above doorways in a home.

5 Experiments in a Natural Home Environment

5.1 Deployment details

Our controlled experiments characterize the sensitivity of height measurements to various conditions, including walking or standing residents, and the effect of shoes. However, these experiments do not reveal the frequency with which these conditions actually occur and affect the measurements in a natural home environment. To evaluate the accuracy of height based identification in such an environment, we deployed ultrasonic sensors in three homes for five days each. The ultrasonic sensors were deployed on doorways of rooms similar to our controlled experiments, and can be seen in figure 4. In addition to the ultrasonic sensors, we deployed the motetrack indoor localization system [13] in all homes to get ground truth locations of residents. Motetrack is a tag and track approach



Fig. 5. Motetrack tags and beacons (left) were used to collect ground truth locations. Motion and magnetic reed switch sensors (middle/right) were used to evaluate STAR.

to localization that requires each resident to carry a mote. It uses trained RSSI signatures from beacon nodes, like the one shown in figure 5a, to localize the mobile motes in the home.

The main goal of the natural home deployments is to evaluate the ability of height sensors to label room visits of residents in the home, *and* compare it to a state of the art non-invasive multi-resident tracking solution that *only* relies on simple activity models derived from labeled binary sensor data [20]. In order to make this comparison, we also deployed anonymous X10 motion sensors in every room, and X10 switch sensors on daily-use objects such as the fridge, microwave, stove etc. Figures 5b and c show examples of the motion and switch sensors used in the homes. In section 6, we also compare with another well studied non-invasive resident identification solution [9] that uses resident mass to differentiate between residents in a home.

Table 1 shows some of the deployment details for the 3 homes, including number of rooms, and ground truth height values of the couple living in each home. Given the large differences in height values in the three homes, we expect our height based identification solution to perform with high accuracy.

5.2 Room occupancy identification in a natural home environment

We evaluate the accuracy with which biometric height sensors are able to identify room visits of residents in a home. We compare the accuracy of our approach with a state of the art passive identification technique based only on 'biometrics' of simple activity models of residents, derived from labeled binary sensor data, as evaluated by Wilson et al in their STAR approach [20].

First, we temporally cluster X10 motion sensor firings from the same room using db-scan [7], to identify discrete room visits of residents in the home; we assume here that these temporal clusters correspond reliably to ground truth room visits of residents. Ground truth resident labels for the temporal clusters are obtained from motetrack's location trace. Our aim is to assign resident labels to each of these clustered room visits, using either biometric height sensors, or using the location trace for each resident computed in STAR using only activity models of residents. To assign resident labels to room visits using only height sensors, we run the log normal MLE classifier on each height event that occurred during the temporal cluster. When the MLE classifier assigns a height event to a resident, that resident is added to the list of labels for that temporal cluster.

The STAR resident tracking system proposed by Wilson et al [20], uses individual Markov state transition and sensor observation models of residents to track their activities and locations. The essence of their tracking approach is that individual residents have different movement/activity patterns in the home, and/or have unique sensor use patterns. Similiar to their original implementation [20], we simply restrict our state space to include current room location of individual residents. The state transition and sensor observation probabilities are learned using counting from our ground truth training data obtained from the motetrack location trace; we performed leave-one-out cross validation over the 5 days of room occupancy data obtained from each home, i.e. for each day, we tracked residents using Markov models trained from all the other days' data. We implemented a multi-hypotheses tracking solution to track room visits of multiple residents in the home, similar to the particle filter solution implemented by Wilson et al [20].

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

Table 1. Details of the 3 homes used in deployments

Figure 5.2 compares the room labeling accuracy of our height sensor approach and the existing approach based on activity and binary sensor use models. We see that identification based on simple activity models and binary sensor use models only achieves accuracy around 65-75% in 3 homes, while height based identification achieves accuracies ranging from 98-100%. Clearly, the activity and *binary* sensor use patterns of residents in these homes are not distinguishing enough to assign room visits with high accuracy. We do not claim here that our approach is better than STAR; instead, we simply compare with an existing instantiation of the STAR framework using only 'biometric' room transition models and binary sensor use models. Certainly, as pointed out by Wilson et al [20], by using more fine-grained sensing at a higher installation cost than our approach, it is possible to better differentiate residents even using these simple models; one could even incorporate height sensor data in the STAR particle filter.

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

cost tag and track solutions [15, 21] may negate the need for such a tracking infrastructure.




Fig. 6. Height sensors achieve higher accuracy than achieved by existing implementations of the STAR approach that use only activity models and binary sensor use models of residents for identification.

6 Accuracy of Height Sensing in Homes Nationwide

We analyze the proportion of homes at a national level where our height solution can *potentially* differentiate residents with high accuracy using the 2006 health and retirement study (HRS 2006) [1], which contains height and weight measures of elderly residents living together in the same household. Of the 2107 multi-resident households with valid height and weight measures for every resident in the home, we used the 2077 households that were two-person households. We do not have currently have access to any anthropometric datasets that support our claim for a wider population demographic. We also note here that wider, longer term deployments in real homes are the best way to evaluate this technology, and our results below are best effort extrapolations from our controlled experiments in the lab.

For each home, using the height values of the residents in the home, we first derive a log normal probability model for each resident in the home using his/her height and the curve fit models described in section 4.2. We then *analytically* calculate the probability that any height event will be assigned to the right resident, assuming that each resident is equally likely to generate a height event. From now on, we refer to this probability as the **probability of correct resident identification** in a home. For each home, HRS 2006 also provides the weight measures of every resident. Jenkins et al [9] in 2007 observe that weight based identification using force plates has a Gaussian error with a mean of 0.67kg and standard deviation of 0.96. Assuming this Gaussian model and mean parameters, we calculate the probability that any gait event will be assigned to the right resident, assuming again that each resident is equally likely to generate a gait event.

Given the probability of correct resident identification for each home in the sample, we compute the proportion of homes where the probability of correct




Fig. 7. Height sensors are potentially applicable with high accuracy to a large proportion of US elderly households. Weight sensors have potentially higher applicability, but require good sensor design to aesthetically install on the floor.

identification is above a fixed threshold; Figure 7 shows how this proportion decreases as we increase the threshold for probability of correct resident identification. Our height based identification solution is potentially applicable to 85% elderly homes in the US with at least 95% identification accuracy. Using force plates and weight based identification, up to 92% of the elderly homes can potentially achieve least 95% identification accuracy. Given the high cost and effort involved in retrofitting a home with force plates, height based identification is preferable, even though it is slightly less accurate; however, alternative cheaper sensing solutions for weight measurement or gait analysis can be explored for preferential use over height sensors in some homes.

7 Improving Height Measurement Accuracy with History

We have shown in the previous sections that height based biometric identification is potentially applicable to a significant proportion of elderly homes in the US with high accuracy. However, from the analysis seen in figure 7, height sensors achieve *less* than 95% identification accuracy in 15% of the homes. In this section, we show how information such as the room topology of a home, and the past history of height sensor events on multiple doorways in a home, can potentially be used to improve the inherent accuracy of biometric identification using the height sensor.

As an illustrative example, assume two residents A and B initially in the bedroom. Assume that after some time, resident A leaves the bedroom, and goes to the kitchen through the living room to get a snack. Even if a few *individual* height sensor events lead to incorrect results from the MLE classifier, the *sequence* of height events generated by A will have a higher likelihood of being classified as resident A; we use spatio-temporal continuity of motion through the constrained floor layout of a home to improve identification accuracy. An assumption in the in the example above, and in the analysis below, is that the

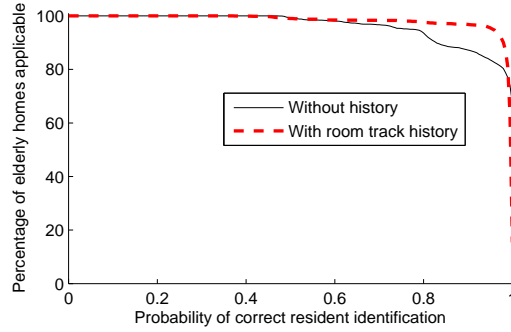


Fig. 8. Our simulation study shows that the history of height measurements collected over the track of a resident through the home potentially improves identification accuracy and applicability of height sensing in US elderly homes.

error at individual height sensors is independent; this assumption may be true most of the time. However, there might occur cases where the error is more systematic, such as a person stooping over to carry a heavy object; in such cases, the utility of using the sequence of height measurements could potentially be reduced.

We use a simulation based study that is driven by the public-use height data from HRS 2006 and our height error models derived in section 4.2, to estimate the improvement in identification accuracy that can be achieved using the history of height events in a home. We assume a 6 room home across all the elderly households for consistency. We define the same HMM for each resident’s room transitions, to indicate equal transition likelihood from one room to another; we do this to ensure that differing room transition patterns of residents or specific room topologies of homes do not unfairly improve the identification accuracy possible by using the past history of height events alone. We generate 1000 height events for each home in HRS 2006, using our HMM to generate room transitions, and using the height error models from section 4.2 to generate noisy height events for each resident. We assume a height sensor at every entry/exit into a room.

Given the simulation trace, we evaluate two approaches to identify resident labels for height events in the home (1) A naive MLE classifier that only considers data from individual height events (2) A probabilistic multi-hypotheses tracker that uses past history and room topology embedded in a HMM. Figure 8 shows the applicability of height based identification across elderly homes in the US with and without history information, based on the results of our simulation experiment. When our probabilistic multi-hypotheses tracker is used, we observe that height based identification can potentially achieve at least 95% identification accuracy in 95% of elderly homes in US, as opposed to only 85% of elderly homes covered by the naive identification algorithm.

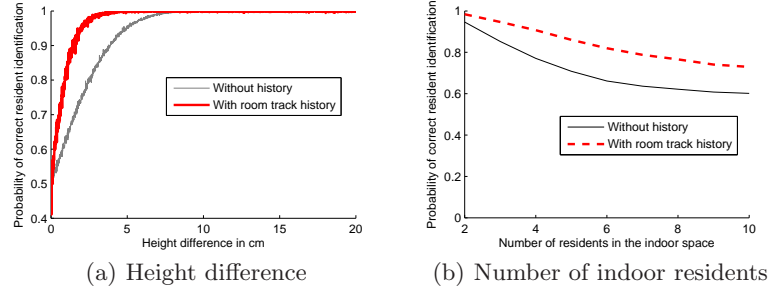


Fig. 9. A history of readings could potentially decrease the height difference required for accurate identification(a), and increase the number of residents that can be reliably differentiated(b)

Figure 9(a) provides more insight into the scenarios where history information might be most useful. When the height difference of the residents living in the home is small, using the past history of height events greatly improves the accuracy over naive MLE classification. Using tracking history can potentially reduce the height difference required for 99% identification accuracy from 6.9cm to 3.25 cm. Figure 9(b) demonstrates another important benefit of using history, the ability to achieve higher accuracy in indoor spaces with more residents; the heights of residents in the hypothetical multi-resident homes (100 sample homes at each point) are randomly generated using height data from the 4154 elderly residents from HRS 2006. By using the history of height events in a home, we can potentially improve the identification accuracy in 4 person homes from 77% to 90%.

8 Discussion

Height measurements can be used for accurate resident room level tracking if height sensors are placed above every entrance and exit to a room, as proposed in this paper. Of course, any biometric sensors including thumbprint or retina scanners [18], could be placed at the entrance of any room to locate residents, but violates our requirement of not requiring manual identification effort from residents. In table 2, we compare the use of height sensors for resident tracking to select existing resident tracking systems in term of four requirements: convenience, deployment time, accuracy, and cost. Deployment time is qualitatively shown for existing solutions we did not implement and is approximated from our empirical deployments in real home environments for solutions we implemented. Table 2 illustrates that, of a representative sample of four existing tracking implementations, none meet all four requirements. Tag and Track systems such as Pedestrian localization [21] and Motetrack [13] can be inconvenient to the user, STAR [20] *using activity models* has low accuracy and requires an inconvenient training phase, and weight sensing *using force plates* or smart floors [9] to track

residents would require a costly installation. Of these existing *implementations*, only our height sensors can meet all four requirements. Of course, exploring alternative implementations of weight sensors using floor mats, or simply including height sensor data in the STAR particle filter, are possible techniques to improve existing *implementations* of these approaches.

Name	Convenience	Deployment time	Accuracy	Cost
Pedestrian Localization	<i>Inconvenient</i>	Very low (5 minutes)	Very high (99-100%)	Very cheap
Motetrack	<i>Inconvenient</i>	Moderate (1 hour)	High, (95-100%)	Affordable
STAR (activity models)	<i>Inconvenient training period (weeks/months)</i>	Moderate (2 hours)	<i>Low, (65-70%)</i>	Affordable
Using Weight (force plates)	Convenient	<i>Very high</i>	High (> 95%)	<i>Very high</i>
Using Height	Convenient	Moderate (1.5 hours)	High (> 95%)	Affordable

Table 2. Evaluation of spectrum of location solutions along four variables - (Convenience, deployment time, Accuracy, Cost) assuming a five room home to deploy in

In this study, we have only explored the use of height sensors above doorways to provide coarse-grained room-level accuracy. Some existing approaches such as tag and track approaches, or invasive camera based approaches, can provide meter-level accuracy; applications that require fine-grained location accuracy would need to install height sensors *inside* rooms, at a higher installation cost. An interesting research question relates to the optimal placement of height sensors inside rooms to help activity inference. A new challenge when using too many height sensors close to each other is multi-path interference affecting the ranging accuracy. Multi-path effects are also an issue in (1) rooms with wide doorways that require several adjacent height sensors to achieve sufficient coverage, and (2) adjacent doorways very close to each other; this needs to be addressed using a distributed synchronization algorithm, or careful placement of the sensors.

In our study, we use a data set which is restricted to elderly residents because it is one of the few public data sets that has both height and weight information for a large number of multi-resident homes. An interesting extension would be to explore how our solution generalizes to a larger population, including young couples, small and large families, and multi-resident student homes, by conducting large scale surveys of anthropometric measures in these homes. Since our

approach is based on a resident biometric, it cannot be applied in all homes with high accuracy. For applications that require higher identification accuracy than offered by height sensing alone, we propose to explore adding multiple non-invasive sensing modalities including floor mats for weight analysis, and color sensors above the doorway. There are also other breakdown scenarios for height sensing that we have not fully explored in this paper. The presence of guests in the home with similar height as the existing residents will reduce identification accuracy. Also, a person who starts to use crutches or a wheelchair, might reduce identification accuracy if her new height corresponds to that of an existing resident in the home.

9 Conclusions

In this work, we demonstrate that ultrasonic range sensors placed above doorways in a home can be used to identify residents with high accuracy as they walk throughout a home, and at the same time satisfy the user requirements of smart home residents. Height is typically a weak biometric, but we make two key insights that make it effective for *in-home* monitoring. First, height is highly effective among small populations where the height differences among residents are likely to be large enough for reliable differentiation. Second, residents walk through the home in predictable, constrained patterns dictated by the floor layout, and the multiple height measurements of the resident as they walk through multiple doorways in the home can be potentially be used to improve the inherent accuracy of the height sensor. In this paper, we quantify the error with which ultrasonic height sensors measure the heights of residents as they walk under the doorway, using both controlled experiments in a lab with 20 subjects, and in 3 real homes. Using publicly available height measures of residents from multi-resident elderly homes and the height error distributions derived from our controlled and in-situ experiments, we extrapolate that a resident identification accuracy of at least 95% can *potentially* be achieved in 85% of elderly homes using a naive classification algorithm and in 95% of elderly homes using our probabilistic multi-hypotheses tracker.

Acknowledgments. This work was supported, in part, by the NSF grant ECCS-0901686 and the NSF grant IIS-0931972. The authors would also like to thank Timothy Hnat from the WSN group at UVa for help with our deployments, and Professor Leo Selavo from the University of Latvia for helpful initial discussions on using ranging for height sensing.

References

1. Health and retirement study 2006. <http://hrsonline.isr.umich.edu>.
2. Quietcare systems - living independently. <http://www.quietcaresystems.com>.
3. Vernier go motion ultrasonic sensor. <http://www.vernier.com/go/gomotion.html>.

4. Wellaware systems for elderly monitoring. <http://www.wellawaresystems.com>.
5. B. Abdelkader et al. Person identification using automatic height and stride estimation. In *International Conference on Pattern Recognition*, 2002.
6. M. Addlesee, A. Jones, F. Livesey, and F. Samaria. The ORL active floor. *IEEE Personal Communications*, 1997.
7. M. Ester et al. A density-based algorithm for discovering clusters in large spatial databases with noise. In *International Conference on Knowledge Discovery and Data Mining*, 1996.
8. G. Gao and K. Whitehouse. The Self-Programming Thermostat: Optimizing Set-back Schedules based on Home Occupancy Patterns. In *First ACM Workshop On Embedded Sensing Systems For Energy-Efficiency In Buildings*, 2009.
9. J. Jenkins and C. Ellis. Using ground reaction forces from gait analysis: body mass as a weak biometric. *Lecture Notes in Computer Science*, 2007.
10. J. Jenkins et al. Weakly identifying system for doorway monitoring. Duke Fontiers Poster Session, May 2006.
11. P. Klasnja, S. Consolvo, T. Choudhury, R. Beckwith, and J. Hightower. Exploring privacy concerns about personal sensing. In *Proceedings of the Seventh International Conference on Pervasive Computing*, Nara, Japan, May 2009.
12. M. K. "ohler, S. Patel, J. Summet, E. Stuntebeck, and G. Abowd. TrackSense: Infrastructure free precise indoor positioning using projected patterns. In *Pervasive computing: 5th international conference, PERVASIVE 2007, Toronto, Canada, May 13-16, 2007: proceedings*, page 334. Springer-Verlag New York Inc, 2007.
13. K. Lorincz and M. Welsh. MoteTrack: a robust, decentralized approach to RF-based location tracking. *Personal and Ubiquitous Computing*, pages 489–503, 2007.
14. Y. Nishida, S. Murakami, T. Hori, and H. Mizoguchi. Minimally privacy-violative human location sensor by ultrasonic radar embedded on ceiling. In *Sensors, 2004. Proceedings of IEEE*, 2004.
15. S. Patel, K. Truong, and G. Abowd. Powerline positioning: A practical sub-room-level indoor location system for domestic use. *Lecture Notes in Computer Science*, 4206:441–458, 2006.
16. G. Shakhnarovich et al. Integrated face and gait recognition from multiple views. In *Conference on Computer Vision and Pattern Recognition*. IEEE Computer Society, 2001.
17. V. Shnayder et al. Sensor networks for medical care. In *Sensys*. ACM, 2005.
18. A. Smith. Exploring the acceptability of biometrics and fingerprint technologies. *International Journal of Services and Standards*, 2005.
19. R. Want, A. Hopper, V. Falcao, and J. Gibbons. The active badge location system. *ACM Transactions on Information Systems (TOIS)*, 1992.
20. D. Wilson and C. Atkeson. Simultaneous tracking and activity recognition (STAR) using many anonymous, binary sensors. In *The Third International Conference on Pervasive Computing*. Springer, 2005.
21. O. Woodman and R. Harle. Pedestrian localisation for indoor environments. In *Proceedings of the 10th international conference on Ubiquitous computing*, pages 114–123. ACM New York, NY, USA, 2008.