

Smart Secure Homes: A Survey of Smart Home Technologies that Sense, Assess, and Respond to Security Threats

Jessamyn Dahmen¹, Diane J. Cook^{1,*}, Xiaobo Wang², and Wang Honglei²

¹School of Electrical Engineering and Computer Science
Washington State University
cook@eecs.wsu.edu

*Corresponding author

²Huawei R&D USA
FutureWei Technologies, Inc.
xiaobo.sc.wang@huawei.com

Abstract

Smart home design has undergone a metamorphosis in recent years. The field has evolved from designing theoretical smart home frameworks and performing scripted tasks in laboratories. Instead, we now find robust smart home technologies that are commonly used by large segments of the population in a variety of settings. Recent smart home applications are focused on activity recognition, health monitoring, and automation. In this paper, we take a look at another important role for smart homes: security. We first explore the numerous ways smart homes can and do provide protection for their residents. Next, we provide a comparative analysis of the alternative tools and research that has been developed for this purpose. We investigate not only existing commercial products that have been introduced but also discuss the numerous research that has been focused on detecting and identifying potential threats. Finally, we close with open challenges and ideas for future research that will keep individuals secure and healthy while in their own homes.

Keywords: security monitoring, anomaly detection, rare event identification, network security, smart home automation

1. Introduction

Over twenty years have passed since Mark Weiser proposed the notion of ubiquitous computing as the new computational model in which technologies disappear into the fabric of everyday life [1]. In many ways, smart homes embody this vision because sensors embedded into everyday environments unobtrusively collect data that monitor the state of the physical environment and its residents while everyday routines are performed. The computational component then reasons about the collected information in order to take an action that optimizes goals such as comfort, safety, or productivity.

While smart homes initially consisted of theoretical designs and smart laboratory experiments, they are rapidly maturing [2]–[5]. The results of this evolution include number of actual prototype smart homes [6]–[13], associated public sensor datasets [14]–[23], and commercial products [24]–[26]. Smart homes are used for a diverse range of applications including activity recognition [27]–[37], health assessment and assistance [38]–[42], environment quality monitoring [43], resource efficiency [44]–[47], and home automation [48].

In this paper, we take an in-depth look at smart home technologies that can be used for home and resident security. Individuals spend a majority of their time in their home or workplace [49], and for many, these places are our sanctuaries. As such, smart home technologies need to contribute to maintaining the safety of residents by preventing as many threats as possible, accurately detecting threats that do occur, and responding quickly and effectively to them.

As Figure 1 shows, a smart home collects data from sensors embedded in the environment. Based on the sensed information, the home reasons about the potential threat and takes an action based on the nature and level of threat that is posed. These three steps – sense, assess, and act - serve as the basis for our coverage of current research and technologies in secure smart homes. We initially describe different types of security issues that smart homes may face and illustrate ways in which the home technology can assist. Second, we then describe current stand-alone sensor systems that detect specific types of threats and summarize current approaches that are taken to responding to threats. Third, we focus on the area that has received the greatest amount of research attention, namely assessing and identifying threats based on sensor data. Finally, we close with a discussion of ongoing challenges for secure smart homes and ideas for future research directions.

Here we motivate the survey of secure smart home technologies through a series of scenarios that illustrate the types of threats that can be encountered in smart homes. Each scenario highlights a different type of security challenge and the role that smart homes can play in assisting with sensing, assessing, and acting on the threat.

Scenario 1: Intruder detection.

Mary lives with her family who are all away for the day. During the morning, the home detects a person's arrival. The house recognizes that the time of day and the type of car fit those of a delivery person. The home provides access to the garage to drop off the items and notifies Mary. When Mary's husband Bob returns home in the afternoon, the home registers his presence and lets Mary know. Late in the evening, the home senses an unusual entry through the window. The camera is turned on to further identify the individual

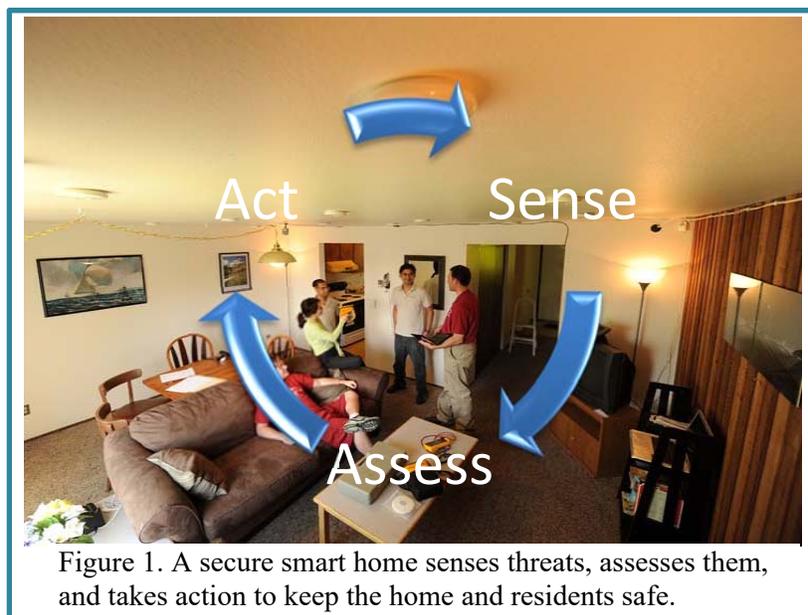


Figure 1. A secure smart home senses threats, assesses them, and takes action to keep the home and residents safe.

and to stream the video to Mary and Bob. They confirm that the individual is their son, who did not have his key and was entering through a window.

Scenario 2: Health event detection. Phil is an 81-year-old man who was diagnosed with Parkinson's Disease five years ago. His mobility has been declining and when getting out of bed one evening he stumbles and falls. Phil is unable to get up to call for help but the home detects the sleep interruption and the subsequent lack of movement. The home asks Phil to confirm he is okay, and when it receives no response, the home contacts emergency services.

Scenario 3: Building system failure detection. Security of smart homes extends beyond individual homes to communities of residents. One complex included fifty apartments, each of which is a smart residence and which share some basic information between them such as indoor air quality, temperature, and electricity usage. When the apartments noted that the levels of volatile organic compounds (VOCs) in five of the residences suddenly rose beyond safe levels, the compound notified the residents in all of the apartments to leave their homes and not return until the situation had been addressed. One of the apartments noted that its resident had been smoking in an apartment which was being remodeled and suggested that the combination of smoke and open toxic chemicals may have contributed to the problem.

These scenarios highlight the diverse nature of security issues that are faced by residents and thus by smart homes as well. A common theme of home security is detection and prevention of intruders, as shown in Scenario 1. However, smart homes that provide security should also be sensitive to health issues that can jeopardize the well-being of residents, as described in Scenario 2. This includes detection of falls, lack of movement, and significant changes in behavioral patterns [50]–[54]. In the same way that the health of a smart home resident can be monitored by a secure smart home system, so the health of the physical home environment can and should be monitored. For example, as described in Scenario 3, the building can be subject to gas leaks, freezing pipes, fires, and other issues that can threaten the health of residents as well as the building [55]. Many of the sensor, assessment, and action strategies can be used across these scenarios as we will see throughout the paper.

While the scenarios illustrate traditional security threats that can be addressed by a smart home, the smart home technology itself can introduce new threats. This motivates the need for smart home systems to be robust and resilient [56]. In particular, if there is a sensor or system failure, the home must still provide protection and needed assistance. Therefore, the smart home itself needs to detect system anomalies and failures in its hardware, software, or communication components [57]. Smart home technologies and the more general Internet of Things technologies also introduce a whole new type of intrusion, namely hacking into the technology infrastructure [58], [59]. Currently smart homes are fairly vulnerable to hacking [60] and this can lead not only to costly pranks (e.g., run the washing machine multiple times) but also life-threatening manipulations (e.g., instead of turning the oven up to 150 degrees, turn the sauna up to 150 degrees) [58].

2. Sensing Threats

As shown in Figure 1, the first step of a secure smart home is to sense the current state of the environment and the residents. Smart home sensors are very diverse and often include a subset of sensors for motion, temperature, lighting, humidity, door use, appliance use, and power consumption, as well as cameras and microphones. With the advent of the Internet of Things (IoT) [61], there is a wealth of devices that provide insights and use the Internet to communicate with each other as well as the resident. In this section, we examine a sampling of technologies that provide sensing capabilities specifically for the purpose of providing a secure environment.

Video cameras are a traditional mechanism for monitoring an environment. They are found in many public venues and provide records of events as well as remote or even automated sensing of threats. Recently, companies including iControl, Nest, SmartThings, Vivint, and Ring have enhanced the traditional camera system for the purpose of smart home safety monitoring. Vivint and Nest cameras can send alerts to homeowners when they detect activity, at which point the resident takes over the task of interpreting the data and acting on it. Ring is unique because it provides a smart doorbell system by

connecting the doorbell to the camera [62]. iControl is even more integrative, because the camera is combined with motion detection, sound detection, and an intruder siren [63]. Alternatively, SmartThings not only facilitates camera-based monitoring and resident alerts, but other devices can be connected as well such as door locks to help residents take remote action in response to possible threats [64].

A second source of ambient sensing in the home for security is audio. Zhuang et al. [65] use Gaussian mixture models to analyze data from a single microphone to specifically detect human falls. Moncrieff et al. [66] scale up the role of the audio signal by quantifying a measure of home “anxiety” based on unusual loud noises that are detected by microphones throughout the home. The microphone is accompanied by a wearable accelerometer to detect whether the resident has experienced a fall.

Commercially-available home security sensing technologies often rely on the resident to interpret data and suggest actions. This process can be made more automated through the use of biometrics. Biometrics will automatically recognize individuals based on unique anatomical traits including voice, gait, retina, and face [67], [68], as well as body shape (anthropometry) [69], footstep shape [70], body weight [71], and heart beat pattern [72]. While biometrics are used frequently for large buildings and operations, they are not as frequently incorporated into individual homes due to the amount of machine learning-based model training that is involved as well as privacy issues. In the context of individual homes, researchers often instead require that residents carry devices to identify themselves [73], [74]. Another approach to recognizing individuals in the home is to recognize behavioral patterns, or behaviometrics [75], rather than just physical properties, or biometrics. Behaviometric-based approaches will be discussed in more detail in Section 4.

An advantage of using the sensor packages described in this section is that they provide a rich source of fine-grained information obtained from video, audio, and specific biometric devices and can ultimately produce a more accurate interpretation of potential threats. This level of information does come at a price, however. Most of the sensors operate with a well-defined field of view, which is the total physical area that is observable by a sensor. Therefore, they need to be placed at locations that would be most likely to encounter the home threats. Employing a large number of such devices would be costly in terms of the initial purchase, the maintenance, the processing of a large amount of data, and the power consumption. On the other hand, utilizing too few devices or placing them in nonoptimal locations will negate their security benefit.

Another challenge posed by these security devices is the potential loss of privacy. Even if the captured information is only stored locally, many individuals feel that the uninterrupted monitoring by the devices is an invasion of their privacy. In fact, many residents turn off these devices when they enter the home [76], relying on the fact that most crimes happen when the home is empty. However, another frequent time for crimes is when the residents are sleeping, and turning off devices in these situations leaves residents vulnerable to threats [77].

3. Acting in Response to Threats

A smart home is typically infused with sensors to monitor the environment. As we described in the last section, these sensors can provide a fairly comprehensive analysis and identification of potential threats. Assuming that the collected information is processed and analyzed for the likelihood and type of threat (discussed in the next section), a smart home will ideally take appropriate steps to act on the threat.

Research and technology development in the area of smart homes has evolved to the point where homes can take autonomous actions in response to detected security or health risks (see Figure 2). As described in Section 2, existing commercial systems automatically provide residents with real-time information when an alert is generated, including notifying them of visitors and providing streaming video identification.

The variety of steps that a smart home can and should take is not limited to alerting and informing the resident, however. In their work, Chitnis et al. [78] surveyed urban, suburban, and rural dwellers from a diversity of backgrounds including homeowners with children who are left unsupervised and individuals with traditional lock-and-key systems. As a result of the survey they proposed an infrastructure that

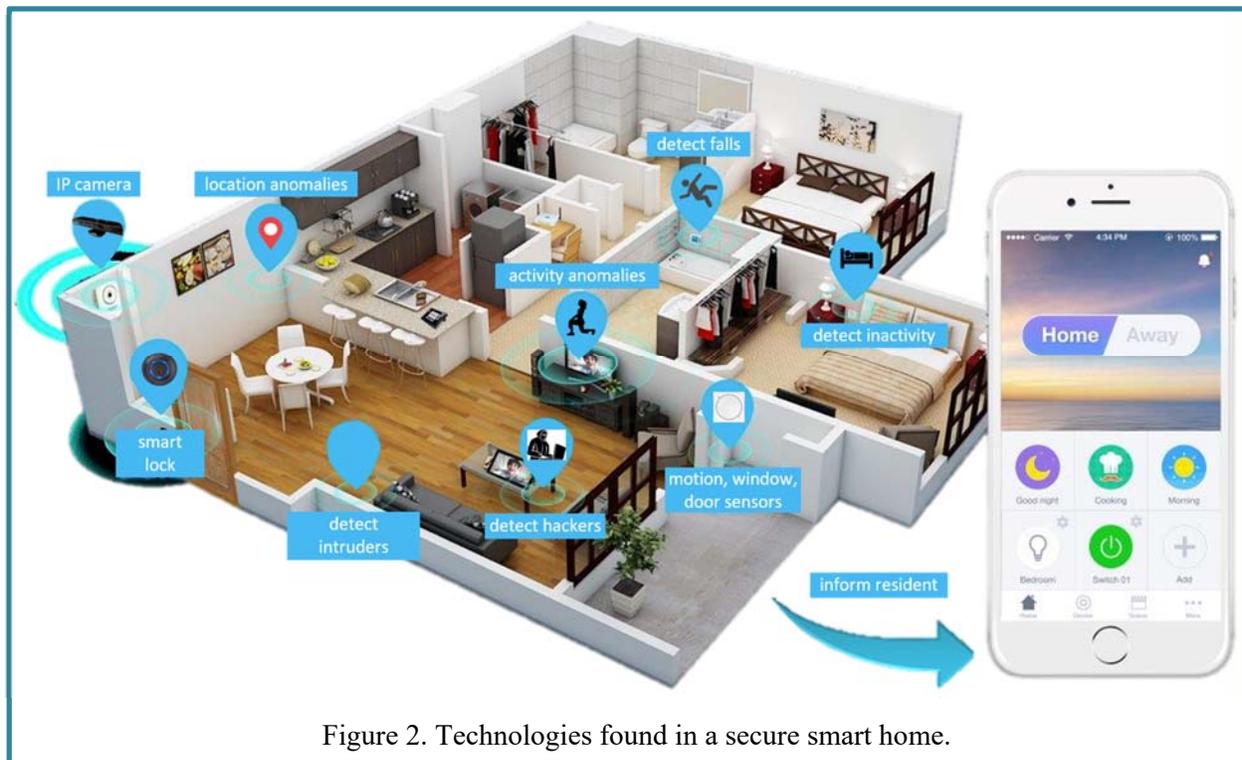


Figure 2. Technologies found in a secure smart home.

granted different types of home access based on biometric matches. As described in Scenario 1 of Section 1, some individuals may only have access to the garage or front porch while repair technicians would also be granted access to areas of the house that need their attention. If an individual manages to enter unauthorized areas of the house, the homeowner is notified.

Homeowners may choose to let ambient sensors run continuously and use the more intensive data-gathering devices such as cameras only when they are out of the home. In such cases, Petersen et al. [79] propose a method to automatically detect these situations and turn on video cameras. In this work, motion and door sensors continuously collect data and a machine learning system is trained to map these sensor readings onto a label indicating whether the residents are at home or away from the home. This approach extracts features including the number of sensor firings during each five-minute interval, an indicator of whether or not the resident is in bed, whether the door sensor was the last reading in the interval, whether the door sensor was the first firing in the interval, and whether the last sensor in the interval emanated from a room connected to an external door. A logistic regressor yielded a sensitivity of 0.939 and a specificity of 0.975 on sample data collected from actual smart homes, which are strong preliminary results supporting this approach.

While intrusion detection is a common application for security systems, much of the technology can also be applied to health monitoring and assistance as well. In the case of work by Dodge et al. [80], by Hodges et al. [81], by Dawadi et al. [82], and by Lotfi et al. [83], unexpected behavioral patterns are viewed as a health risk for individuals who are at risk of cognitive decline. These researchers have found that an increase in the number of activity anomalies and variation in behavior patterns such as activity times and walking speed are correlated with changes in cognitive health. As in the case with the intrusion detection research, these findings provide insights that can be used by smart homes in order to keep residents safe. For example, residents and their caregivers can use this information to change the level of care that the individual needs.

In research by Ali et al. [84] and of Das et al. [85], threats are detected in the form of abnormalities in how residents perform their daily activities. For many individuals, these variations would not be considered a risk. However, for individuals with memory limitations, performing daily activities independently is critical. Functional impairment has been associated with increased health care use and

placement in long-term care facilities [86], [87], days in the hospital [88], falls [89], conversion to dementia [90], [91], and morbidity and mortality [92]. When an abnormality is detected, the individual can be prompted for the next activity step to help them keep on track and successfully complete the activity without caregiver intervention. This in turn increases functional independence and reduces the burden for caregivers.

4. Detecting and Assessing Threats

In this section, we close the loop shown in Figure 1. Both research and commercial efforts have made contributions in the areas of developing sensors for secure homes and acting autonomously or in partnership with residents to respond to threats. The largest body of research, however, has focused on the middle step, analyzing collected sensor data to detect and assess potential security threats. We organize our discussion of threat assessment in order of scale. We start with describing approaches to detect specific security-related situations, move toward summarizing approaches that perform general detection of threat-based anomalies, and finish with a discussion of security-based research in other fields that can impact future work on secure smart homes.

4.1. Detecting resident-based target states

As demonstrated by the scenarios at the beginning of this paper, security is a broad term that encompasses many aspects of health and safety. Instead of trying to tackle the problem of home-based security as a whole, some researchers focus on one piece of the problem. One targeted situation is one in which the home resident is inactive for an unusually long period of time. Cuddihy et al. [93] focus on detecting times with unusual bouts of inactivity, which can occur when a resident falls or has other health issues that may require intervention. This is a situation that is of particular interest for smart home developers but is difficult to address without special-purpose hardware. In this approach, Cuddihy employs motion and door sensors in smart homes that are trained based on 10 days of data to learn normal ranges of activity / inactivity for each 30-minute interval throughout the day. The main challenge with software-based approaches to detecting inactivity is a high rate of false positives. As a result, their evaluation on smart home data focused on reducing the number of alerts to a goal of at most one per month (indicative of an acceptable number of false positives), which they satisfied for 91% of the collected data.

Another situation that has received a great deal of attention is detecting when an individual falls in their home [51], [52], [94]. Approximately one in every three older adults experiences a fall [95]. Because the global population is aging and individuals want to stay in their own homes, detecting such falls is a concern not only for older adults and their caregivers but also for society as a whole. While research on using wearable sensors to detect falls based on accelerometer and gyroscope readings abounds [51], [96], detecting falls with home-based ambient sensors can be more difficult. Some smart home-based research has been pursued, however. These methods typically use microphones [97] or video data [98] to determine when a fall has occurred.

Of these approaches, video methods have been shown to be the most able to discern falls because of the fine-grained information that is available. However, video approaches also raise concerns for maintaining resident privacy. This concern can be addressed in part by using depth cameras to capture only figure silhouettes rather than detailed images. Stone and Skubic [94] introduce such a method using the Microsoft Kinect. Theirs is a two-stage method that first extracts the foreground information and second generates a confidence that a fall has occurred. This method is trained and tested on 3,339 days of continuous data for 16 residents in independent apartments containing 454 falls (9 of which naturally occurred during the data collection). To highlight the foreground information, the background is subtracted from the Kinect data by maintaining a distribution of pixel values as shown in Figure 3. Assuming that background information is less dynamic, depth camera data is subtracted if it matches the earlier distributions, indicating that it represents “inactive” information. The resulting extracted information is the dynamic foreground information. Once features are extracted from the foreground

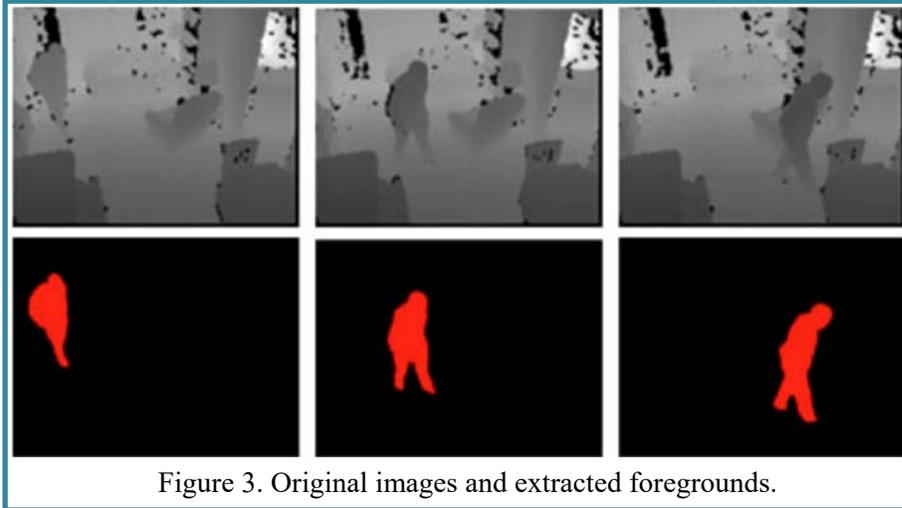


Figure 3. Original images and extracted foregrounds.

depth information, an ensemble of decision stumps (single-feature decision trees) is used to generate a confidence that a fall has occurred.

Another approach that does not rely on video data is to use vibration sensors in the floor to detect falls. As an example, Alwan et al. [99] embed a piezoelectric sensor in the floor. Vibration patterns are learned for typical

movements such as walking and used to distinguish typical movements from possible falls. The challenge for all of these systems is obtaining enough training data to detect and recognize the many types of falls that can occur such as slips, trips, stumbles, and slumps.

Audio sensors face similar benefits and concerns. Microphone arrays can provide fine-grained information to sense the state of the home. With the increasing popularity of voice-interactive products such Amazon Echo or Google Home, they will be easy to integrate into lifestyles and serve multiple purposes. As with other sources of rich information, however, they do provide a privacy challenge that will need to be addressed in ongoing research.

4.2. Detecting home-based target states

In this previous section, we described methods that detect target states for a smart home resident that can represent health or security risks. In contrast, here we examine approaches to detecting specific home states that highlight security needs. These approaches target recognizing whether individuals are in the home and whether the individuals that are present are residents, visitors, or intruders. The first method we include is a machine learning method that Teoh and Tan [76] designed specifically to recognize intruders. The authors intend homeowners to use this method as a precaution to maintain security when they are out of the home. A neural network is trained based on input from motion sensors, closed-circuit televisions, RFID tags, magnetic contact switches, and glass breakage sensors. The neural network maps extracted features from these sensors onto a binary output indicating whether the current scenario is an intrusion or not. As can be imagined, the most difficult challenges with this approach are obtaining realistic training data (without inviting criminals to break into test homes!) and dealing with the severe class imbalance that will result.

The next step in recognizing anomalous home behavior is to model movement patterns through the home with the purpose of determining whether or not visitors are present. This goal is complementary to the work of Teoh and Tan because it operates effectively when the resident is at home. In this case, the algorithm can be used to inform the resident that others are in the home. If the visitors are unexpected then this can represent a risk not only to the home and belongings but to the resident as well. In this work, Aicha et al. [100] model transitions between sensors that are typical when the resident is alone and when visitors are present. For a particular time slice t , $N^M(t)$ represents the number of transitions between motion sensors that are not near either other in the smart home and $N^D(t)$ represents transitions that include the front door sensor. The approach employs a Markov Modulated Poisson Process, or MMPP. Because the MMPP is built on a Markov chain it follows expected sequences between resident alone time, visitor time, and unusual absence of visitors, as shown in Figure 4. In addition, the MMPP utilizes a Poisson process that can model periodic influences such as daily visits from a caregiver, weekly visits

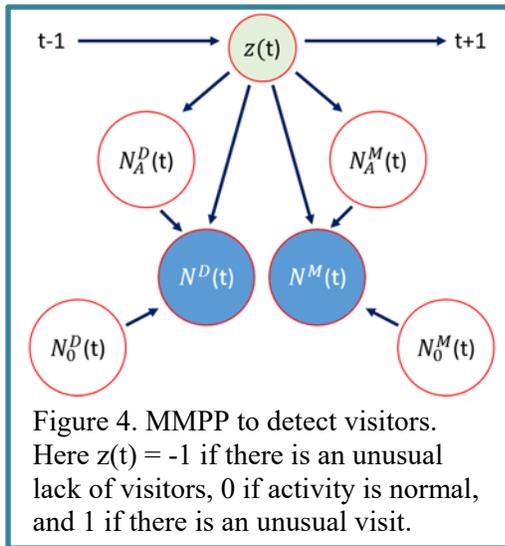
from a maid service, and monthly visits from children. The model outputs a value, $z(t)$, that indicates whether the current activity in the home is normal, if there is an unusual visit, or if there is an unusual lack of visits. An appealing feature of this approach is that it can scale to larger numbers of residents and visitors. In contrast to work by Petersen et al. that also detect whether visitors are in the home [101], the approach described by Aicha is unsupervised. As a result, it does not require that large amounts of labeled training data are available. The approach was tested on continuous data collected in actual smart homes for 16 residents. The data was based on documented visits from cleaners, care providers, and family members.

The final work we include in this section considers the final piece of the home state puzzle, namely whether the resident is at home. Petersen et al. [79] seek to determine whether the resident is in or out of the home in order to assess loneliness and well-being of older adults. However, this is valuable for security purposes because times when the resident is out of the home are particularly at risk for intruders or for dangerous changes in the environment (e.g., gas leaks, water leaks) that may need to receive a timely intervention. These researchers collected data from smart home sensors and extracted relevant features for each five-minute time window including number of sensor firings, presence in bed, and front door usage. The smart homes included a video camera that was used for ground truth labels. Instead of collecting continuous video data, video was only recorded for 5 seconds each time motion was detected. A logistic regressor yielded a sensitivity of 0.393 and specificity of 0.975 for data collected from 150 older adults for 30 days each.

4.3. Detecting home-based anomalies

Targeting specific resident states or house states is valuable for home-based security because the targeted states are known to be security risks. While recognizing these states is challenging, they are simpler than the task of finding a more varied class of security threats using a single technique. To find a larger collection of possible threats, researchers often turn to anomaly detection. As Chandola states, anomaly detection is “the problem of finding patterns in data that do not conform to expected behavior” [102]. We first focus on specific classes of anomalies that can be detected directly with smart home sensors then transition in the next section to finding anomalies based on activities that a resident performs in the home. There are several standard techniques for finding anomalies or outliers that are commonly employed. One such technique is to cluster, or group, data points into clusters based on their distance to the cluster center. Once the data is clustered, points that are far from all of the cluster centers can be labeled as outliers and can be considered as anomalies. In cases where the data is normally distributed, z scores can also be computed. The z score for any data point is its distance from the sample mean, divided by the sample standard deviation. Z scores greater than 3.5 are typically considered to be outliers. Although the methods in this section do not rely on activity recognition, they are still more complex than is found with many anomaly and outlier detection methods. This is because such methods are typically univariate and thus cannot handle the multi-variable, complex data collected by smart home sensors.

In our treatment of home-based anomalies, we start at the lowest information level and work our way up to more abstract information. As described by Youngblood and Cook [103], the lowest level of a smart home infrastructure is the physical sensors. Correspondingly, there are research groups who identify anomalies based on unusual sensor readings and unusual sequences of sensor firings. As an example, Ordonez et al. [104] use low-level sensor data to identify outliers indicative of behavioral anomalies. This group collects data of typical sensor timings for motion, pressure mat, and door switch sensors. From this data, Bayesian statistics are used to capture the typical sensor firing time, the sensor firing sequence, and the duration of each sensor state. Outliers can be detected based on a Bayesian model and used to indicate anomalies.



Haque et al. [105] recognize the fact that some sensor-based outliers may be due to sensor faults rather than changes in the system being monitored. To mitigate the false alarms that can be generated due to sensor failures, they dynamically adjust decision thresholds. Traditional SMO-based regression is used to generate a predicted sensor reading. If the difference between predicted and actual values is greater than a threshold than an anomaly is reported, and the threshold is adjusted incrementally to reflect the actual amount of error that is normal for a given individual. Majority voting is then used to combine anomaly scores for all of the sensors and an anomaly is reported only if the majority votes for this label. The combination of threshold adjusting and majority voting reduces the occurrence of false alarms that can happen due to issues such as failure of individual sensors.

In a smart home, the location of residents within a home can be fairly easily triangulated using infrared motion sensors that fire whenever a heat mass the size of a small child or larger moves in its field of regard. Door sensors, window sensors, and object sensors also help to identify where people are located inside the home. There is thus a natural progression toward monitoring time spent in locations throughout the home in order to determine anomalies. Lotfi et al. [83] take this approach to analyzing smart home data collected in the homes of older adults. Each firing of a motion or door sensor is represented as a tuple consisting of the sensor location, the sensor firing time, and amount of time the sensor remains in the “ON” or “OPEN” state, indicating that the resident is staying in one location. They cluster all of the data points and use the techniques described earlier such as distance from cluster centers to identify points that are potential anomalies in terms of a resident’s location. These points can be used to inform a care provider to check on the health of the resident.

Aran et al. [106] take a similar approach to the one proposed by Lotfi et al. Like Lotfi’s work, Aran infers the current location of a smart home resident and uses k means clustering to group similar data points based on time of day, location, and time spent in the location. The clustering algorithm provides a convenient basis for determining outliers as data points that do not fit well in any of the existing clusters. Unlike the Lotfi study, Aran includes not only locations in the home but also time spent out of the home. Anywhere outside the sensed home is treated as yet another location and can be monitored for unusual times the resident is out of the home, unusual-length stays in the home without going out, and unusually long or short outings.

Novák et al. [107] also represent data points as a combination of sensor location, time, and duration in a particular state. Like Lotfi, the goal of this group is to provide early alerts for anomalous behavior based on unusual location-based activity in terms of being in an unusual place at an unusual time, staying in the location an unusually long period of time, or staying in the location an unusually short period of time. Instead of applying a k means clustering algorithm, Novák uses the collected data to learn a Self Organization Map (SOM), which is a type of self-structuring neural network. The SOM itself visualizes points that are far from its neighbors and can therefore be further analyzed as potential anomalies.

Virone et al. [108] integrate their anomaly detection algorithm into a complete graphical alert system for residents and caregivers that provides a visual indication of behavioral outliers at different time scales. This group observes the amount of time that is spent in a room of the house on each occasion as a separate data point. Data is broken into single-hour time windows and these points produce circadian rhythms over a twenty four-hour period that can be analyzed for normalcy and anomalies. As with the other projects, this group detects unusually long or short durations spent in a room and also detects unusually high or low levels of activity within the room at a given time. Figure 5 shows the interface displaying various alarms, where each activity is associated with a particular location in the house.

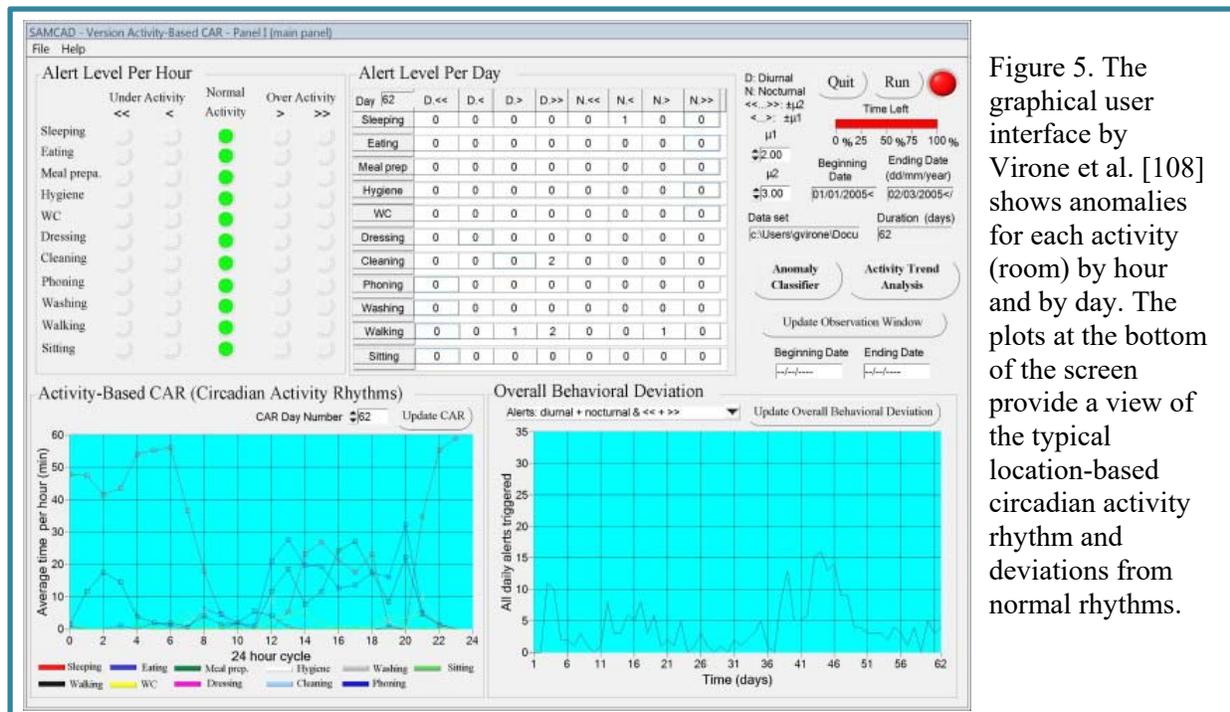


Figure 5. The graphical user interface by Virone et al. [108] shows anomalies for each activity (room) by hour and by day. The plots at the bottom of the screen provide a view of the typical location-based circadian activity rhythm and deviations from normal rhythms.

The papers summarized in this section focus on locations where a smart home resident spends time. This focus simplifies the data analysis to an extent because the current activity does not need to be inferred. Some of the papers categorized as location-based anomaly detection do refer to monitoring activities but they equate activity with a distinct region of the home (or outside the home). For some activities this approach will work, but in other cases a single room may be used for multiple activity functions (e.g., the bedroom may be used for reading, working, and watching TV in addition to sleeping). In the next section, we consider approaches that look for anomalies at the level of individual activities.

4.4. Detecting activity anomalies

We now look at detection of anomalies at a deeper level of detail, namely within and among activities that are tracked in a smart home. Learning and understanding observed activities is at the center of many fields of study. The challenge of activity recognition is to automate activity learning by mapping smart home sensor data to a label that indicates the corresponding activity that an individual is performing. Activity recognition data consists of sensor firings with the corresponding sensor identifier and time. Features are extracted from sensor data at a particular time and a supervised learning algorithm maps the features onto a value from a list of possible activities (e.g., sleep, eat, cook, take medicine, exercise, work, read) which indicates the activity that is currently being performed. Because of the insight that automated activity recognition sheds on human behavior and the valuable context activity labels bring to smart homes, activity recognition is a highly-investigated area of research [27], [29], [30], [109], [110]. Activity models are created from a variety of methods including support vector machines, Gaussian mixture models, decision trees, and probabilistic graphs.

As in earlier sections, we first examine approaches that investigate specific constrained situations then transition to consider more general approaches. An example that falls into the first category is work by Han et al. [111]. This approach assumes that an activity recognizer is available that labels sensor data with activity labels for eating, toileting, and sleeping. They also monitor overall movement in the home and participant weight. Han uses recognized activities to identify specific changes in behavior. If the participant is diagnosed as depressed they look for changes of concern that include less movement around the house, decrease in hygiene and eating, and more severe sleep disturbances. If the participant is

diagnosed as diabetic they look for more frequent eating, drinking, sleeping, and toileting. These specific changes are known to be problematic for the target populations and represent a health risk when they are recognized by a smart home.

Williams and Cook [112] also use activity recognition to look for specific classes of changes. They are particularly interested in using smart homes to perform health monitoring for individuals who may be experiencing sleep disturbances due to PTSD or decline in cognitive health. After labeling sensor data with forty possible activity labels, they segment the data into waketime behavior and sleeptime behavior, then collapse all of the activity-based behavior parameters into two scores, one for wake and one for sleep. They then use past and current wake scores to forecast the upcoming sleep value and use past and current sleep scores to forecast the next day's wake value. Their motivation is to predict and detect sleep disturbances in order to circumvent them by suggesting changes in wake behavior that may prevent the predicted sleep problems.

Analysis of the literature reveals a number of approaches that apply activity recognition to label sensor data, then look for anomalies in the sequences of activities that are observed. These methods are very similar to those of the location-based anomaly detection approaches described earlier in this section, although they are applied to activity sequences rather than location sequences. Work by Mocanu and Florea [113], Cardinaux et al. [114], Elbert et al. [115], and Mori et al. [116] all develop automated approaches to detecting deviations in daily activity patterns. These include activity start times, activity durations, activity locations, and activity orders. These methods employ different models of activities but all share a motivation of finding changes in activity routines that indicate a health or security risk. In all cases the method is tested on a very small set of constrained activities but have the potential to be scaled to a richer set of activity details.

All anomaly detection algorithms suffer from a common problem: detecting too many false positive anomalies. Without constraining the type of anomalies that are of interest, most methods find excessive numbers of outliers. Many of these can be due to harmless situations such as visitors, trips away from home, and changes in smart home hardware. Hoque et al. [117], [118] address this challenge by integrating expert-provided explanations for known situations that may appear anomalous. In their approach, this group uses hierarchical clustering to group activity data points. Expert-provided rules are integrated for known events that may look like anomalies to an automated system. This approach could be further extended to continually obtain expert feedback that explains detected anomalies so that the number of false positives decreases and the number of true positives increases over time, not just for the detected anomalous situations but for others that are similar.

The activity-based anomaly methods that have been described up to this point model anywhere from 3 to 40 activities. In everyday behavior, though, these predefined activities only comprise at most 50% of an individual's daily routine [29]. In order to be thorough, smart home security systems need to model all aspects of a normal routine in order to catch every possible anomaly and important behavior change. Wang et al. [119] partially address this situation by discovering frequent sensor firing sequences that represent "activities", then detecting anomalies among occurrences of these sequences. This is a promising direction for research but is extremely dependent upon the discovery method. This particular approach focuses purely on sensor state duration and resident trajectory, which may be insufficient for representing and analyzing complex activities.

Once activities are modeled and understood, anomalies can then be found within each activity. These anomalies may indicate that an intruder is attempting to emulate a known activity (and is making some mistakes) or that the resident is experiencing difficulty in remembering how to successfully complete an activity. A secure smart home can detect such a situation and alert the resident in order to ensure that the home is safe and to assist the resident if needed in completing a task. Tong et al. [120] tackle this problem using hidden state conditional random fields (HCRF). A HCRF can be used to compare an observed activity with a database of activities performed normally. The most similar activity from the database provides a basis of comparison and the degree of dissimilarity between the observed activity and its closest match indicates the "abnormality" of the observed behavior.

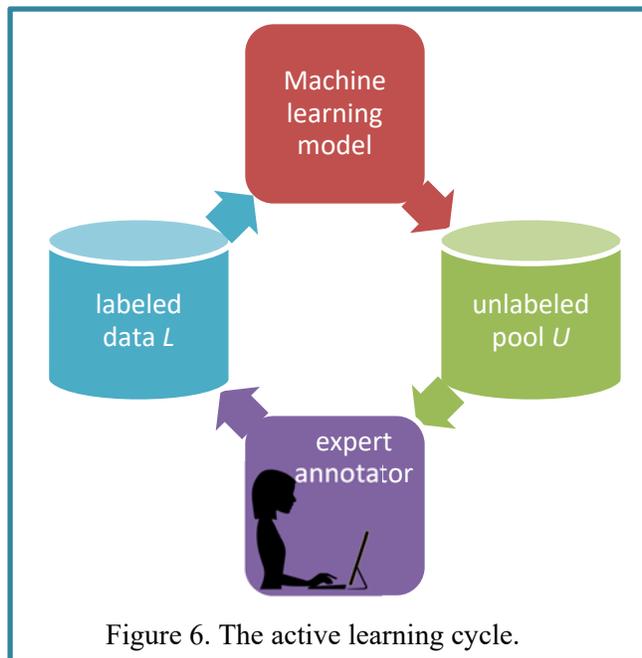


Figure 6. The active learning cycle.

Das et al. [85] take this a step further by decomposing each modeled activity into individual steps. A one-class classifier is used to determine whether an activity occurrence is “normal”. If it is not, the type of error can be identified as an error of omission (a step is missing), substitution (an incorrect tool is used), irrelevant action (unnecessary steps are included), or inefficient action (a step is included that slows down or compromises the efficiency of action completion). This method can be used to not only determine if activity steps are unusual and therefore suspect, but also the type of error that was committed in order to identify a security risk or to correct the incomplete activity.

The paper that we have discussed come directly from the literature on behavior-based anomaly detection. However, techniques have been introduced in related fields that can be employed for activity-based anomaly detection

and maintaining secure smart homes. As an example, the field of active learning for activity modeling is focused on detecting unusual activity-based data points. In this case, however, the reason for detecting these points is to determine the data points that will most greatly benefit from expert guidance in terms of providing an explanation and label for the data point. Active learning employs a machine learning algorithm to interactively query an expert (for example, a human annotator) to obtain a label for a sensor data point pulled from a pool of unlabeled data, U . Once the label is provided the data point is added to the labeled training set, L , and used to update the activity models for improved recognition performance, as shown in Figure 6. Active learning techniques often select data points that exhibit properties such as the most labeling uncertainty [121], [122] or the least consensus among a committee of classifiers [123]. As an example, the Optimized Probabilistic Active Learning (OPAL) [124] algorithm picks data points for labels that are not well understood by the current activity models and which are outliers with respect to the entire set of data points. These methods work closely with existing anomaly detection algorithms and activity classifiers and therefore can be useful as a component of behavior-based anomaly detection.

The last topic we consider in this section is rare event detection [125]–[128]. Rare events are events that occur very infrequently. While they may not be considered as anomalies by methods described in this section, they are found in less than 10% of the data and when they do occur the consequences are often negative and can be dramatic. This is a topic that is often investigated in areas such as insurance risk modeling, web modeling, and hardware fault detection. In many cases, researchers have found that normal events are similar to each other and rare events are quite different not only from normal events but from each other. For example, in the arena of credit card transactions normal transactions are very standard but fraudulent use varies in many ways. Anomaly detection can be used in some cases for rare event detection. However, the nature of rare events may be better understood than the general class of anomalies. For example, intrusion into a home is rare and negative but is understood. Similarly, a fall in the home is understood and while it may therefore not be considered an anomaly, it is important to detect. For this reason, supervised machine learning techniques can be used but researchers must design specialized methods to deal with the extreme class imbalance that occurs in these cases.

As we conclude the discussion of threat detection, we point out that much of the existing research in the field has focused on detecting behavior-based anomalies for health applications (such as the one described in Scenario 2 at the beginning of the paper) rather than for home security. Monitoring the health of a smart home resident is indeed a critical part of ensuring that the resident is safe and secure in his or her own home. However, these techniques can be easily adapted to for intrusion detection and building-

based anomalies, as well. Changes in a well-established routine may indicate that the person in the home is not the resident or that the building is not operating as normal. In any of these situations, steps should be taken by the home and/or the resident to make sure they are protected.

4.5. Security research in related areas

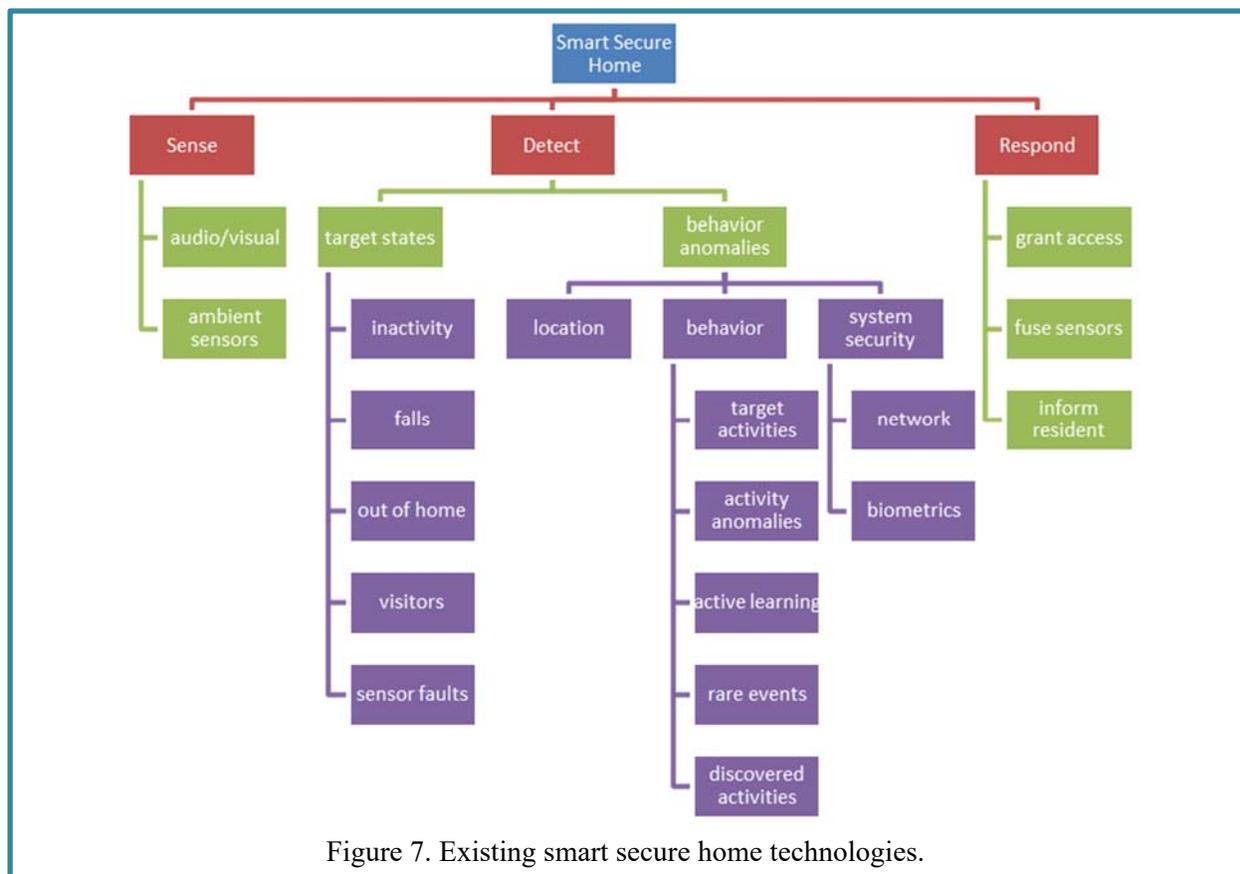
Much of the discussion in this paper has encompassed papers published on the topic of smart homes, behavior analysis, and building automation. To be thorough, we also need to highlight work in related areas that can be valuable when creating smart secure homes. Sensor data collected in smart homes is temporal by nature and anomaly detection techniques for time series data are abundant. These techniques often employ statistical methods like a moving average, an autoregressive moving average (ARMA), or an autoregressive integrated moving average (ARIMA) to predict a future value of a variable given past values. If the observed value is unexpected given the mean and variance of the prediction then an anomaly is detected. Kalman filters can extend these basic ideas to combine multiple sources of information when the information is possibly noisy. Change point detection methods [129] can also be applied to time series data such as smart home sensor firings. In this case, data from two consecutive time periods are examined to determine if they come from the same probability distribution. If the data are sufficiently different then a change has occurred, possibly due to a transition to a new activity or due to an anomaly.

There are also a number of approaches to anomaly detection in graph-based data [130]–[133]. Graphs provide a natural representation for data that is rich in structure because independence is not assumed between individual data points. Instead, the nodes of a graph can represent the individual data points, or features, and nodes are connected by labeled or unlabeled edges when a relationship is exhibited between the points. While these methods employ some of the same basic ideas for anomaly detection as we have described for activity analysis and time series analysis, additional features can be extracted and integrated into the analysis which have been shown to be particularly valuable for security applications [134]–[136].

Many of the time series-based anomaly detection methods quantify the degree of surprise that a data point exhibits and if the value is greater than a threshold, an anomaly is reported. As we have discussed throughout this paper, though, the methods frequently generate a large number of false positive anomalies. In the time series literature, randomization tests [137] are employed to reduce the number of false positives and this technique may be valuable for smart home applications as well. In this method, a generative model such as a hidden Markov model is learned from actual data and used to create synthetic data that is similar to the real data. A large number (e.g., $n=1000$) of anomaly detection runs can be performed on the simulated data. The fraction of runs in which the surprise value of the data exceeded the real data surprise value can then be used as a p-value for the real anomaly and provides a method of quantifying confidence in the detected anomaly.

We have focused up to this point on services that a smart home can provide in terms of preventing, detecting, and responding to security threats. A natural issue arises when the smart home capabilities themselves are compromised. In such a case, the smart home can become a liability for a resident rather than an asset. Cases in which smart homes and smart home components are hacked are becoming well documented. Some investigators such as Hill [60] and Fernandes et al. [138] actually broke into and controlled other homes in order to assess the security of smart home infrastructures. Others have virtually entered private homes for nefarious reasons and have taken over the home's web cameras or devices [139], [140]. Recent efforts [141] have demonstrated that smart lighting technologies are susceptible to infiltration and even smart door locks themselves are vulnerable to a whole new breed of burglar: the computer hacker [142].

Smart home hackers often enter the premises by gaining access to a network [143]. However, relying on stand-alone security technologies does not ensure that they cannot be broken. For example, spoofing techniques can be used to imitate (spoof) picture or video of a resident in order to fool a camera [144] and can be accomplished on portable devices for ease of use by intruders [145]. Research such as the work by



Lai and Tai [146] on detecting biometric spoofing thus needs to regularly be pursued to detect and prevent such attempts.

The ability to hack into connected systems is by no means a new issue that researchers face [58], [147]. Computer networks regularly face password attackers, sniffers that spy on network user traffic, IP spoofers, and man-in-the-middle attackers by individuals to intercept and alter communication between people or devices. Attackers may seek to steal information, gain access to devices, corrupt data, or introduce malicious data and programs. Because of the increasing maturity of smart homes, however, smart home researchers, designers, and residents now need to be aware of these issues. Just as we have demonstrated in this paper for smart home data, anomaly detection is a popular topic for network researchers as well [148] to detect and prevent these threats [143]. Smart homes offer unprecedented means to provide safety, security, comfort and productivity. Continued research is necessary to ensure that the security benefits from these homes outweigh the potential risks.

5. Conclusions and Future Research Directions

In this paper, we examined the potential and existing technologies that can transform any home into a smart secure home. As the summary in Figure 7 shows, a wide range of individual technology components have been developed to sense, detect, assess, and respond to a variety of secure threats in home settings. As our summary has made clear, there is also a tremendous need for ongoing research to improve these technologies and to prevent new security risks from arising as a result of transforming homes into smart homes.

This look at existing work raises issues that can be addressed with new research efforts. Throughout the discussion we highlighted the fact that it is difficult to characterize what type of anomaly will be of interest for smart home residents. In fact, the distinguishing factor between data outliers, behavioral

anomalies, and rare events of interest is the explanation for the event. Continued research can thus help to formalize the nature of behavioral anomalies that are particularly critical for security application. This step will allow data mining techniques to focus more effectively on detecting appropriate patterns in smart home data.

Another direction for continued research is to automate the explanation of anomalous events. For example, Sprint et al. [54] employ a virtual classifier to distinguish normal from abnormal behavior in a way that generates rule-based explanations for the differences. This approach could be used to more quickly verify the accuracy and criticality of the detected anomaly. It would also be interesting to see if anomaly explanations can generalize to entire populations. For example, what may first seem like an anomaly may indeed be a fall that occurred in the home and is thus a rare event. This type of event can then be learned for all smart homes, lowering false positive rates for new smart homes and reducing the amount of efforts residents must expend in explaining the events and training the system.

Yet another consideration for future research and development is sensor fusion. Instead of relying on a single activity-based anomaly detector or biometric device, information can be combined from multiple sources to increase confidence in the assessment of collected information. They can also complement each other's strengths and weaknesses. For example, when behavior analysis detects a potential anomaly, other components such as a camera system can be initiated to confirm to deny that the situation is threatening.

Finally, ongoing research must consider human factors when designing smart secure homes. Research would be valuable to consider how many false positives residents will tolerate as well as what type of system training they can provide. The actions that a home takes to ensure resident safety must be understandable and comfortable for residents or users will quickly remove all smart components from their home. Once the user understands the workings of the system and comes to know what to expect in its actions, he or she can trust that the home will truly act in a way to keep them secure and comfortable.

Acknowledgements

This work was supported in part by the National Science Foundation under grant 121407.

References

- [1] M. Weiser, "The computer for the Twenty-First Century," *Sci. Am.*, vol. 165, pp. 94–104, 1991.
- [2] C. Wilson, T. Hargreaves, and R. Hauxwell-Baldwin, "Smart homes and their users: A systematic analysis and key challenges," *Pers. Ubiquitous Comput.*, vol. 19, no. 2, pp. 463–476, 2015.
- [3] M. R. Alam, M. B. I. Reaz, and M. A. M. Ali, "A review of smart homes - past, present, and future," *IEEE Trans. Syst. Man, Cybern. Part C*, vol. 42, no. 6, pp. 1190–1203, 2012.
- [4] T. Jones, "Artificial intelligence coming to a home near you," *Digital Construction*, 2012. [Online]. Available: <http://www.constructiondigital.com/innovations/artificial-intelligence-coming-to-a-home-near-you>.
- [5] T. Cohen, "I'm afraid I can't let you do that, Dave': Scientists predict 'smart' homes controlled by computer will be a reality in 10 years," *Mail Online*, 2012. [Online]. Available: <http://www.dailymail.co.uk/sciencetech/article-2122343/Scientists-predict-smart-homes-controlled-reality-10-years.html>.
- [6] G. Abowd and E. D. Mynatt, "Designing for the human experience in smart environments," in *Smart Environments: Technologies, Protocols and Applications*, 2005, pp. 153–174.
- [7] H. Hagraas, F. Doctor, A. Lopez, and V. Callaghan, "An incremental adaptive life long learning approach for type-2 fuzzy embedded agents in ambient intelligent environments," *IEEE Trans. Fuzzy Syst.*, vol. 15, no. 1, pp. 41–55, 2007.

- [8] S. Intille, K. Larson, E. Munguia-Tapia, J. Beaudin, P. Kaushik, J. Nawyn, and R. Rockinson, "Using a live-in laboratory for ubiquitous computing research," in *Pervasive*, 2006, pp. 349–365.
- [9] A. Helal, W. Mann, H. Elzabadiani, J. King, Y. Kaddourah, E. Jansen, H. El-Zabadani, and Y. Kaddoura, "The Gator Tech Smart House: A programmable pervasive space," *IEEE Comput.*, vol. 38, no. 3, pp. 50–60, Mar. 2005.
- [10] M. C. Mozer, "Lessons from an adaptive home," in *Smart Environments: Technology, Protocols, and Applications*, D. J. Cook and S. K. Das, Eds. Wiley, 2004, pp. 273–298.
- [11] D. J. Cook, M. Youngblood, E. Heierman, K. Gopalratnam, S. Rao, A. Litvin, and F. Khawaja, "MavHome: An agent-based smart home," in *Pervasive Computing*, 2003, pp. 521–524.
- [12] D. J. Cook, A. Crandall, B. Thomas, and N. Krishnan, "CASAS: A smart home in a box," *IEEE Comput.*, vol. 46, no. 7, pp. 62–69, 2012.
- [13] Philips., "365 days' ambient intelligence research in HomeLab." 2003.
- [14] S. Intille, J. Nawyn, B. Logan, and G. Abowd, "Developing shared home behavior datasets to advance HCI and ubiquitous computing research," in *International Conference on Human Factors in Computing Systems Extended Abstracts*, 2009, pp. 4763–4766.
- [15] ASU, "Sensor activity prediction in smart homes," 2012. .
- [16] Boxlab, "List of home datasets," 2012. [Online]. Available: <https://boxlab.wikispaces.com/List+of+Home+Datasets>.
- [17] F. De la Torre, J. Hodgins, J. Montano, S. Valcarcel, and J. Macey, "Guide to the Carnegie Mellon University multimodal activity (CMU-MMAC) database," 2009.
- [18] E. Kim, S. Helal, J. Lee, and S. Hossain, "The making of a dataset for smart spaces," in *International Conference on Ubiquitous Intelligence and Computing*, 2011.
- [19] J. Malik, S. Petrov, A. Berg, and S. Petrox, "Action recognition datasets," 2012. [Online]. Available: <http://www.eecs.berkeley.edu/Research/Projects/CS/vision/action/>.
- [20] C. Wren, Y. Ivanov, D. Leigh, and J. Westhues, "The MERL motion detector dataset," in *Workshop on Massive Datasets*, 2007, pp. 10–14.
- [21] A. Bulling, U. Blanke, and B. Schiele, "A tutorial on human activity recognition using body-worn inertial sensors," *ACM Comput. Surv.*, vol. 46, no. 3, pp. 107–140, 2014.
- [22] CASAS, "WSU CASAS Datasets," 2016. [Online]. Available: <http://ailab.wsu.edu/casas/datasets/>.
- [23] University of Florida, "Ambient intelligence datasets," <http://www.cise.ufl.edu/~prashidi/Datasets/ambientIntelligence.html>. .
- [24] Samsung SmartThings, "Stay connected to your home and family," 2016. [Online]. Available: <https://www.smarthings.com/>.
- [25] Honeywell, "Your connected home," 2016. [Online]. Available: http://homesecurity.honeywell.com/home_automation.html.
- [26] Google, "Get to know Google Home," 2016. [Online]. Available: <https://madeby.google.com/home/>.
- [27] L. Chen, J. Hoey, C. D. Nugent, D. J. Cook, and Z. Yu, "Sensor-based activity recognition," *IEEE Trans. Syst. Man, Cybern. Part C Appl. Rev.*, vol. 42, no. 6, pp. 790–808, 2012.
- [28] N. Krishnan and D. J. Cook, "Activity recognition on streaming sensor data," *Pervasive Mob. Comput.*, vol. 10, pp. 138–154, 2014.
- [29] D. J. Cook and N. Krishnan, *Activity Learning: Discovering, Recognizing, and Predicting Human Behavior from Sensor Data*. Wiley, New York, 2015.
- [30] J. K. Aggarwal and M. S. Ryoo, "Human activity analysis: A review," *ACM Comput. Surv.*, vol. 43, no. 3, pp. 1–47, 2011.

- [31] L. Chen and I. Khalil, "Activity recognition: Approaches, practices and trends," in *Activity Recognition in Pervasive Intelligent Environments*, L. Chen, C. D. Nugent, J. Biswas, and J. Hoey, Eds. Atlantis Ambient and Pervasive Intelligence, 2011, pp. 1–31.
- [32] P. Tuaraga, R. Chellappa, V. S. Subrahmanian, O. Udrea, and P. Turaga, "Machine recognition of human activities: A survey," *IEEE Trans. Circuits Syst. Video Technol.*, vol. 18, no. 11, pp. 1473–1488, 2008.
- [33] I. L. Liao, D. Fox, and H. Kautz, "Location-based activity recognition using relational Markov networks," in *International Joint Conference on Artificial Intelligence*, 2005, pp. 773–778.
- [34] E. Munguia-Tapia, S. S. Intille, and K. Larson, "Activity recognition in the home using simple and ubiquitous sensors," in *Pervasive*, 2004, pp. 158–175.
- [35] H. Fang and C. Hu, "Recognizing human activity in smart home using deep learning algorithm," in *Chinese Control Conference*, 2014, pp. 4716–4720.
- [36] P. Roy, S. Giroux, B. Bouchard, A. Bouzouane, C. Phua, A. Tolstikov, and J. Biswas, "A possibilistic approach for activity recognition in smart homes for cognitive assistance to Alzheimer's patients," *Atl. Ambient Pervasive Intell.*, vol. 4, pp. 33–58, 2011.
- [37] A. Fleury, N. Noury, and M. Vacher, "Supervised classification of activities of daily living in health smart homes using SVM," in *Proceedings of the International Conference of the IEEE Engineering in Medicine and Biology Society*, 2009, pp. 6099–6102.
- [38] P. Dawadi, D. J. Cook, and M. Schmitter-Edgecombe, "Automated clinical assessment from smart home-based behavior data," *IEEE J. Biomed. Heal. Informatics*, 2016.
- [39] D. J. Cook, P. Dawadi, and M. Schmitter-Edgecombe, "Analyzing activity behavior and movement in a naturalistic environment using smart home techniques," *IEEE J. Biomed. Heal. Informatics*, vol. 19, no. 6, pp. 1882–1892, 2015.
- [40] M. E. Morris, B. Adair, O. Miller, R. Hansen, A. Pearce, N. Santamaria, L. Viegas, M. Long, and C. Said, "Smart home technologies to assist older people to live well at home," *J. Aging Sci.*, vol. 1, pp. 1–9, 2013.
- [41] L. Walsh, A. Kealy, J. Loane, and J. Doyle, "Inferring health metrics from ambient smart home data," *IEEE Int. Conf. Bioinforma. Biomed.*, 2014.
- [42] J. Hoey, A. Monk, and A. Mihailidis, "People, sensors, decisions: Customizable and adaptive technologies for assistance in healthcare," *ACM Trans. Interact. Intell. Syst.*, vol. 2, no. 4, 2012.
- [43] S. Deleawe, J. Kuszniir, B. Lamb, and D. J. Cook, "Predicting air quality in smart environments," *J. Ambient Intell. Smart Environ.*, vol. 2, no. 2, pp. 145–154, 2010.
- [44] Y. Riche, J. Dodge, and R. Metoyer, "Studying always-on electricity feedback in the home," in *International Conference on Human Factors in Computing Systems*, 2010, pp. 1995–1998.
- [45] I. B. P. P. Dinata and B. Hardian, "Predicting smart home lighting behavior from sensors and user input using very fast decision tree with Kernel Density Estimation and improved Laplace correction," in *International Conference on Advanced Computer Science and Information Systems*, 2014, pp. 171–175.
- [46] A. Fensel, S. Tomic, V. Kumar, M. Stefanovic, S. V. Aleshin, and D. O. Novikov, "SESAME-S: Semantic smart home system for energy efficiency," *Informatik-Spektrum*, vol. 36, no. 1, pp. 46–57, 2013.
- [47] S. Gupta, M. S. Reynolds, and S. N. Patel, "ElectriSense: Single-point sensing using EMI for electrical event detection and classification in the home," in *ACM International Conference on Ubiquitous Computing*, 2010, pp. 139–148.
- [48] J. Scott, A. J. B. Brush, J. Krumm, B. Meyers, M. Hazas, S. Hodges, and N. Villar, "PreHeat: Controlling home heating using occupancy prediction," in *International Conference on Ubiquitous Computing*, 2011, pp. 281–290.

- [49] Bureau of Labor Statistics, “American time use survey,” 2016. [Online]. Available: <http://www.bls.gov/tus/>.
- [50] M. Skubic, B. H. Harris, E. Stone, K. C. Ho, B.-Y. Su, and M. Rantz, “Testing non-wearable fall detection methods in the homes of older adults,” in *IEEE International Conference of the Engineering in Medicine and Biology Society*, 2016, pp. 557–560.
- [51] M. Mubashir, L. Shao, and L. Seed, “A survey on fall detection: Principles and approaches,” *Neurocomputing*, vol. 100, pp. 144–152, 2013.
- [52] N. Noury, T. Herve, V. Rialle, G. Virone, E. Mercier, G. Morey, A. Moro, and T. Porcheron, “Monitoring behavior in home using a smart fall sensor and position sensors,” in *International Conference on Microtechnologies in Medicine and Biology*, 2000, pp. 607–610.
- [53] G. Sprint and D. J. Cook, “Unsupervised detection and analysis of changes in everyday physical activity data,” *J. Biomed. Inform.*, 2016.
- [54] G. Sprint, D. J. Cook, R. Fritz, and M. Schmitter-Edgecombe, “Using smart homes to detect and analyze health events,” *IEEE Comput.*, 2016.
- [55] G. Demiris and B. K. Hensel, “Technologies for an aging society: A systematic review of ‘smart home’ applications,” *IMIA Yearb. Med. Informatics*, vol. 47, no. 1, pp. 33–40, 2010.
- [56] S. Guillet, B. Bouchard, and A. Bousouane, “Correct by construction security approach to design fault tolerant smart homes for disabled people,” in *International Conference on Emerging Ubiquitous Systems and Pervasive*, 2013, pp. 257–264.
- [57] E. Pardo, D. Espes, and P. Le-Parc, “A framework for anomaly diagnosis in smart homes based on ontology,” in *International Conference on Ambient Systems, Networks and Technologies*, 2016, pp. 545–552.
- [58] N. Komninos, E. Philippou, and A. Pitsillides, “Survey in smart grid and smart home security: Issues, challenges, and countermeasures,” *IEEE Commun. Surv. Tutorials*, vol. 16, no. 4, pp. 1933–1954, 2014.
- [59] D. Storm, “Of 10 IoT-connected home security systems tested, 100% are full of security FAIL,” *computerworld.com*, 2015. .
- [60] K. Hill, “When ‘smart homes’ get hacked: I haunted a complete stranger’s house via the Internet,” *Forbes*, 2013. [Online]. Available: <http://www.forbes.com/sites/kashmirhill/2013/07/26/smart-homes-hack/>.
- [61] E. Brown, “Who needs the internet of things?,” *linux.com*, 2016. .
- [62] Ring, “Never miss a visitor. With Ring, you’re always home.,” 2016. [Online]. Available: <https://ring.com/>.
- [63] Icontrol Networks, “Home security,” 2016. [Online]. Available: <https://getpiper.com/howitworks/>.
- [64] SmartThings, “Discovery ways to use SmartThings for monitoring and security,” 2016. [Online]. Available: <https://www.smarthings.com/uses/monitoring-security>.
- [65] X. Zhuang, J. Huang, G. Potamianos, and M. Hasegawa-Johnson, “Acoustic fall detection using Gaussian mixture models and GMM supervectors,” in *IEEE International Conference on Acoustics, Speech, and Signal Processing*, 2009, pp. 69–72.
- [66] S. Moncrieff, S. Venkatesh, G. West, and S. Greenhill, “Multi-modal emotive computing in a smart house environment,” *Pervasive Mob. Comput.*, vol. 3, no. 2, pp. 79–94, 2007.
- [67] A. K. Jain and K. Nandakumar, “Biometric authentication: System security and user privacy,” *IEEE Comput.*, vol. 45, no. 11, pp. 87–92, 2012.
- [68] euronews, “Smarter home security camera recognises intruders says maker,” 2016. [Online]. Available: <http://www.euronews.com/2016/08/03/smarter-home-security-camera-recognises-intruders-says-maker>.

- [69] V. Andersson, R. Dutra, and R. Araujo, "Anthropometric and human gait identification using skeleton data from Kinect sensor," in *ACM Symposium on Applied Computing*, 2014, pp. 60–61.
- [70] A. Helal, W. Mann, H. Elzabadi, J. King, Y. Kaddourah, and E. Jansen, "Gator Tech Smart House: A Programmable pervasive space," *IEEE Comput. Mag.*, pp. 64–74, Mar. 2005.
- [71] J. Jenkins and C. Ellis, "Using ground reaction forces from gait analysis: Body mass as a weak biometric," in *Pervasive Computing*, 2007, pp. 251–267.
- [72] K. Watanabe, Y. Kurihara, and H. Tanaka, "Ubiquitous health monitoring at home - sensing of human biosignals on flooring, on tatami mat, in the bathtub, and in the lavatory," *IEEE Sens. J.*, vol. 9, no. 12, pp. 1847–1855, 2009.
- [73] N. Matsushita, S. Tajima, Y. Ayatsuka, and J. Rekimoto, "Wearable key: Device for personalizing nearby environment," in *International Symposium on Wearable Computers*, 2000, pp. 119–126.
- [74] A. Venkatesh, "Digital home technologies and transformation of households," *Inf. Syst. Front.*, vol. 10, no. 4, pp. 391–395, 2008.
- [75] A. Crandall and D. J. Cook, "Behaviometrics for multiple residents in a smart environment," in *Human Aspects in Ambient Intelligence*, 2013, pp. 55–71.
- [76] C. Teoh and C. Tan, "A neural network approach towards reinforcing smart home security," in *Asia-Pacific Symposium on Information and Telecommunication Technologies*, 2010.
- [77] F. Cardinaux, S. Brownsell, M. Hawley, D. Bradley, S. Chitnis, N. Deshpande, A. Shaligram, B. Das, D. J. Cook, N. Krishnan, M. Schmitter-Edgecombe, M. Hodges, N. Kirsch, M. Newman, M. Pollack, E. Hoque, R. Dickerson, S. Preum, M. Hanson, A. Barth, J. Stankovic, N. Komninos, E. Philippou, A. Pitsillides, R. J. Robles, T. Kim, J. Williams, D. J. Cook, M. Alam, N. Roy, M. Petruska, A. Zemp, H. Ali, D. G. Amalarethinam, D. T. Anderson, M. Ros, J. M. Keller, M. P. Cuellar, M. Popescu, M. Delgado, A. Vila, O. Aran, D. Sanchez-Cortes, M. T. Do, D. Gatica-Perez, I. Batal, D. Fradkin, J. Harrison, F. Moerchen, M. Hauskrecht, V. Chandola, A. Banerjee, V. Kumar, J.-J. Chen, Z.-X. Jiang, Y.-L. Chen, W.-T. Wu, J.-M. Liang, G. Civitarese, C. Bettini, S. Belfiore, P. Cuddihy, J. Weisenberg, C. Graichen, M. Ganesh, G. Demiris, B. K. Hensel, W. Eberle, L. Holder, B. Massengill, D. Elbert, H. Storf, M. Eisenbarth, O. Unalan, M. Schmitt, S. Guillet, B. Bouchard, A. Bousouane, M. Gupta, J. Goa, C. Aggarwal, J. Han, Y. Han, M. Han, S. Lee, A. M. J. Sarkar, Y.-K. Lee, S. Haque, M. Rahman, A. Aziz, D. Harrison, W. Seah, R. Rayudu, E. Hoque, J. Stankovic, A. Lazarevic, J. Srivastava, V. Kumar, A. Lotfi, L. C., S. M. Mahmoud, M. J. Akhlaghinia, A. Lotfi, C. Langensiepen, I. Mocanu, A. M. Florea, T. Mori, A. Fujii, M. Shimosaka, H. Noguchi, T. Sato, N. Noury, T. Herve, V. Rialle, G. Virone, E. Mercier, G. Morey, A. Moro, T. Porcheron, M. Novak, J.F., L.L., F. Ordonez, P. de Toldeo, A. Sanchis, E. Pardo, D. Espes, P. Le-Parc, T. Senator, H. Goldberg, A. Memory, C. Teoh, C. Tan, Y. Tong, R. Chen, J. Gao, G. Virone, C. Wang, Q. Zheng, Y. Peng, D. De, W.-Z. Song, P. Wang, K.-M. Chao, C.-C. Lo, W.-H. Lin, H.-C. Lin, W.-J. Chao, M. Xie, S. Han, B. Tian, and S. Parvin, "An investigative study for smart home security: Issues, challenges and countermeasures," *Sensors*, vol. 15, no. 4, pp. 1–15, 2016.
- [78] S. Chitnis, N. Deshpande, and A. Shaligram, "An investigative study for smart home security: Issues, challenges and countermeasures," *Wirel. Sens. Netw.*, vol. 8, pp. 61–68, 2016.
- [79] J. Petersen, D. Austin, J. A. Kaye, M. Pavel, and T. L. Hayes, "Unobtrusive in-home detection of time spent out-of-home with applications to loneliness and physical activity," *IEEE J. Biomed. Heal. Informatics*, vol. 18, no. 5, pp. 1590–1596, 2014.
- [80] H. H. Dodge, N. C. Mattek, D. Austin, T. L. Hayes, and J. A. Kaye, "In-home walking speeds and variability trajectories associated with mild cognitive impairment," *Neurology*, vol. 78, no. 24, pp. 1946–1952, 2012.
- [81] M. Hodges, N. Kirsch, M. Newman, and M. Pollack, "Automatic assessment of cognitive impairment through electronic observation of object usage," in *International Conference on*

- Pervasive Computing*, 2010, pp. 192–209.
- [82] P. Dawadi, D. Cook, and M. Schmitter-Edgecombe, “Modeling patterns of activities using activity curves,” *Pervasive Mob. Comput.*, 2015.
- [83] A. Lotfi, L. C., S. M. Mahmoud, and M. J. Akhlaghinia, “Smart homes for the elderly dementia sufferers: identification and prediction of abnormal behavior,” *J. Ambient Intell. Humaniz. Comput.*, vol. 3, pp. 205–218, 2012.
- [84] H. Ali and D. G. Amalarethinam, “Detecting abnormality in activities performed by people with dementia in smart environment,” *Int. J. Comput. Sci. Inf. Technol.*, vol. 5, pp. 2453–2457, 2014.
- [85] B. Das, D. J. Cook, N. Krishnan, and M. Schmitter-Edgecombe, “One-class classification-based real-time activity error detection in smart homes,” *IEEE J. Sel. Top. Signal Process.*, 2016.
- [86] D. Marson and K. Hebert, “Functional assessment,” in *Geriatric Neuropsychology Assessment and Intervention*, 2006, pp. 158–189.
- [87] A. Desai, G. Grossberg, and D. Sheth, “Activities of daily living in patients with dementia: Clinical relevance, methods of assessment and effects of treatment,” *CNS Drugs*, vol. 18, pp. 853–875, 2004.
- [88] U. Sonn, G. Grimbyand, and A. Svanborg, “Activities of daily living studied longitudinally between 70 and 76 years of age,” *Disabil. Rehabil.*, vol. 18, pp. 91–100, 1996.
- [89] S. Zimmerman and J. Magaziner, “Methodological issues in measuring the functional status of cognitively impaired nursing home residents: the use of proxies and performance-based measures,” *Alzheimer Dis. Assoc. Disord.*, vol. 8, pp. S281–S290, 1995.
- [90] P. Barberger-Gateau, J. Dartigues, and L. Letenneur, “Four instrumental activities of daily living score as a predictor of one-year incident dementia,” *Age Ageing*, vol. 22, pp. 457–463, 1993.
- [91] K. Peres, V. Chrysostome, C. Fabrigoule, J. Orgogozo, J. Dartigues, and P. Barberger-Gateau, “Restriction in complex activities of daily living in MCI,” *Neurology*, vol. 67, pp. 461–466, 2006.
- [92] F. Nourhashemi, S. Andrieu, S. Gillette-Guyonnet, B. Vellas, J. Albarede, and H. Grandjean, “Instrumental activities of daily living as a potential marker of frailty: a study of 7364 community-dwelling elderly women (the EPIDOS study),” *J. Gerontechnology*, vol. 56A, pp. M448–M453, 2001.
- [93] P. Cuddihy, J. Weisenberg, C. Graichen, and M. Ganesh, “Algorithm to Automatically Detect Abnormally Long Periods of Inactivity in a Home,” in *ACM SIGMOBILE International Workshop on Systems and Networking Support for Healthcare and Assisted Living Environments*, 2007, pp. 89–94.
- [94] E. Stone and M. Skubic, “Fall detection in homes of older adults using the Microsoft Kinect,” *IEEE J. Biomed. Heal. Informatics*, vol. 19, no. 1, pp. 290–301, 2015.
- [95] S. R. Lord, C. Sherrington, and H. B. Menz, “Falls in older people: Risk factors and strategies for prevention,” Cambridge, England, 2001.
- [96] A. K. Bourke, J. Klenk, L. Schwickert, K. Aminian, E. A. F. Ihlen, S. Mellone, J. L. Helbostad, L. Chiari, and C. Becker, “Fall detection algorithms for real-world falls harvested from lumbar sensors in the elderly population: A machine learning approach,” in *IEEE Annual International Conference of the Engineering in Medicine and Biology Society*, 2016, pp. 1–6.
- [97] Y. Li, L. Zeng, M. Popescu, and K. C. Ho, “Acoustic fall detection using a circular microphone array,” in *IEEE Annual International Conference of the Engineering in Medicine and Biology Society*, 2010, pp. 2242–2245.
- [98] C. Rougier, J. Meunier, A. St-Arnaud, and J. Rousseau, “Robust video surveillance for fall detection based on human shape deformation,” *IEEE Trans. Circuits Syst. Video Technol.*, vol. 21, no. 5, pp. 611–622, 2011.

- [99] M. Alwan, P. J. Rajendran, S. Kell, D. Mack, S. Dalal, M. Wolfe, and R. Felder, "A smart and passive floor-vibration based fall detector for elderly," in *IEEE International Conference on Information and Communication Technology*, 2006, pp. 1003–1007.
- [100] A. N. Aicha, G. Englebienne, and B. Krose, "Modeling visit behaviour in smart homes using unsupervised learning," in *ACM Conference on Ubiquitous Computing*, 2014, pp. 1193–1200.
- [101] J. Petersen, N. Larimer, J. A. Kaye, M. Pavel, and T. L. Hayes, "SVM to detect the presence of visitors in a smart home environment," in *International Conference of the IEEE Engineering in Medicine and Biology Society*, 2012, pp. 5850–5853.
- [102] V. Chandola, A. Banerjee, and V. Kumar, "Anomaly detection: A survey," *ACM Comput. Surv.*, vol. 41, pp. 1–15, 2009.
- [103] G. M. Youngblood and D. J. Cook, "Data mining for hierarchical model creation," *IEEE Trans. Syst. Man, Cybern. Part C*, vol. 37, no. 4, pp. 1–12, 2007.
- [104] F. Ordonez, P. de Toldeo, and A. Sanchis, "Sensor-based Bayesian detection of anomalous living patterns in a home setting," *Pers. Ubiquitous Comput.*, vol. 19, pp. 259–270, 2015.
- [105] S. Haque, M. Rahman, and A. Aziz, "Sensor anomaly detection in wireless sensor networks for healthcare," *Sensors*, vol. 15, pp. 8764–8786, 2015.
- [106] O. Aran, D. Sanchez-Cortes, M. T. Do, and D. Gatica-Perez, "Anomaly detection in elderly daily behavior in ambient sensing environments," in *Human Behavior Understanding*, 2016, pp. 51–67.
- [107] M. Novak, J.F., and L.L., "Anomaly detection in user daily patterns in smart-home environment," *J. Sel. Areas Heal. Informatics*, vol. 3, pp. 1–11, 2013.
- [108] G. Virone, "Assesing everyday life behavioral rythms for the older generation," *Pervasive Mob. Comput.*, vol. 5, pp. 606–622, 2009.
- [109] T. Barger, D. Brown, and M. Alwan, "Health status monitoring through analysis of behavioral patterns," *IEEE Trans. Syst. Man, Cybern. Part A*, vol. 35, no. 1, pp. 22–27, 2005.
- [110] S.-R. Ke, H. L. U. Thuc, Y.-J. Lee, J.-N. Hwang, J.-H. Yoo, and K.-H. Choi, "A review on video-based human activity recognition," *Computers*, vol. 2, no. 2, pp. 88–131, 2013.
- [111] Y. Han, M. Han, S. Lee, A. M. J. Sarkar, and Y.-K. Lee, "A framework for supervising lifestyle diseases using long-term activity monitoring," *Sensors*, vol. 12, pp. 5363–5379, 2012.
- [112] J. Williams and D. Cook, "Forecasting behavior in smart homes based on past sleep and wake patterns," *Technol. Heal. Care*, 2016.
- [113] I. Mocanu and A. M. Florea, "A model for activity recognition and emergency detection in smart environments," in *International Conference on Ambient Computing, Applications, Services and Technologies*, 2011, pp. 13–19.
- [114] F. Cardinaux, S. Brownsell, M. Hawley, and D. Bradley, "Modelling of behavioural patterns for abnormality detection in the context of lifestyle reassurance," *Prog. Pattern Recognition, Image Anal. Appl.*, vol. 5197, pp. 243–251, 2008.
- [115] D. Elbert, H. Storf, M. Eisenbarth, O. Unalan, and M. Schmitt, "An approach for detecting deviations in daily routine for long-term behavior analysis," in *In Pervasive Health*, 2011, pp. 426–433.
- [116] T. Mori, A. Fujii, M. Shimosaka, H. Noguchi, and T. Sato, "Typical behavior patterns extraction and anomaly detection algorithm based on accumulated home sensor data," in *Conference on Future Generation Communication and Networking*, 2007.
- [117] E. Hoque, R. Dickerson, S. Preum, M. Hanson, A. Barth, and J. Stankovic, "Holmes: A comprehensive anomaly detection system for daily in-home activities," in *International Conference on Distributed Computing in Sensor Systems*, 2015, pp. 40–51.
- [118] E. Hoque and J. Stankovic, "Semantic anomaly detection in daily activities integrate expert rules

- for acceptable anomalies,” in *ACM International Joint Conference on Pervasive and Ubiquitous Computing*, 2012, pp. 633–634.
- [119] G. Okeyo, L. Chen, H. Wang, and R. Sterritt, “Dynamic sensor data segmentation for real-time knowledge-driven activity recognition,” *Pervasive Mob. Comput.*, vol. 10, pp. 155–172, 2014.
- [120] Y. Tong, R. Chen, and J. Gao, “Hidden state conditional random field for abnormal activity recognition in smart homes,” *Entropy*, vol. 17, pp. 1358–1378, 2015.
- [121] M. Dredze and K. Crammer, “Active learning with confidence,” in *Proceedings of ACL*, 2008, pp. 233–236.
- [122] A. J. Joshi, F. Porikli, and N. Papanikolopoulos, “Multi-class active learning for image classification,” in *IEEE Conference on Computer Vision and Pattern Recognition*, 2009.
- [123] Y. Freund and R. E. Schapire, “Experiments with a new boosting algorithm,” in *International Conference on Machine Learning*, 1996, pp. 148–156.
- [124] G. Kreml, D. Kottke, and V. Lemaire, “Optimised probabilistic active learning (OPAL) for fast, non-myopic, cost-sensitive active classification,” *Mach. Learn.*, vol. 100, no. 2, pp. 449–476, 2015.
- [125] A. Lazarevic, J. Srivastava, and V. Kumar, “Data Mining for Analysis of Rare Events: A Case Study in Security, Financial and Medical Applications,” in *Pacific-Asia Conference on Knowledge Discovery and Data Mining*, 2004.
- [126] D. Harrison, W. Seah, and R. Rayudu, “Rare event detection and propagation in wireless sensor networks,” *ACM Comput. Surv.*, vol. 48, p. 58, 2016.
- [127] D. Pelleg and A. W. Moore, “Active learning for anomaly and rare-category detection,” in *Advances in Neural Information Processing Systems*, 2004, pp. 1073–1080.
- [128] S. Koh and S. D. Ravana, “Unsupervised rare pattern mining: A survey,” *ACM Trans. Knowl. Discov. Data*, vol. 10, no. 4, p. 45, 2016.
- [129] S. Aminikhanghahi and D. J. Cook, “A survey of methods for time series change point detection,” *Knowl. Inf. Syst.*, pp. 1–29, Sep. 2016.
- [130] L. Akoglu, H. Tong, and D. Koutra, “Graph based anomaly detection and description: A survey,” *Data Min. Knowl. Discov.*, vol. 29, no. 3, pp. 626–688, 2015.
- [131] C. Noble and D. J. Cook, “Graph-based anomaly detection,” in *Proceedings of the Ninth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2003.
- [132] W. Eberle, L. Holder, and B. Massengill, “Graph-based anomaly detection applied to homeland security cargo screening,” in *Florida Artificial Intelligence Research Society Conference*, 2012.
- [133] S. Rayana and L. Akoglu, “Less is more: Building selective anomaly ensembles,” *ACM Trans. Knowl. Discov. Data*, vol. 10, no. 4, p. 42, 2016.
- [134] W. Eberle and L. Holder, “Scalable anomaly detection in graphs,” *Intell. Data Anal.*, vol. 19, pp. 57–74, 2015.
- [135] D. Cook, L. Holder, S. Thompson, P. Whitney, and L. Chilton, “Graph-based analysis of nuclear smuggling data,” *J. Appl. Secur. Res.*, vol. 4, no. 4, pp. 501–517, 2009.
- [136] D. Chakrabarti, Y. Zhan, D. Blandford, C. Faloutsos, and G. Blelloch, “NetMine: New mining tools for large graphs,” in *SIAM Workshop on Link Analysis, Counter-terrorism and Privacy*, 2004.
- [137] B. Efron and R. J. Tibshirani, *An introduction to the bootstrap*. CRC Press, 1994.
- [138] E. Fernandes, J. Jung, and A. Prakash, “Security analysis of emerging smart home applications,” in *IEEE Symposium on Security and Privacy*, 2016, pp. 636–654.
- [139] A. Lee, “Hacking the connected home: When your house watches you,” *readwrite*. [Online]. Available: <http://readwrite.com/2013/11/13/hacking-the-connected-home-when-your-house->

watches-you#feed=/tag/connected-home&awesm=~osmDA6o9bkgx84.

- [140] T. Clemons, “Wake up call: Mom learns daughters’ bedroom webcam was hacked,” 2016. .
- [141] C. O’Flynn, “A lightbulb worm?,” 2016. .
- [142] A. Rose and B. Ramsey, “Picking Bluetooth low energy locks from a quarter mile away,” *DefCon*, 2016. .
- [143] P. Wang, K.-M. Chao, C.-C. Lo, W.-H. Lin, H.-C. Lin, and W.-J. Chao, “Using malware for software-defined networking-based smart home security management through a taint checking approach,” *Int. J. Distrib. Sens. Networks*, vol. 12, no. 8, p. 2016, 2016.
- [144] A. Hadid, “Face biometrics under spoofing attacks: Vulnerabilities, countermeasures, open issues and research directions,” in *IEEE Conference on Computer Vision and Pattern Recognition Workshops*, 2014, pp. 113–118.
- [145] Y. Xu, T. Price, J.-M. Frahm, and F. Monrose, “Virtual U: Defeating face liveness detection by building virtual models from your public photos,” in *USENIX Security Symposium*, 2016, pp. 497–512.
- [146] C. Lai and C. Tai, “A smart spoofing face detector by display features analysis,” *Sensors*, vol. 16, no. 7, pp. 1136–1150, 2016.
- [147] R. J. Robles and T. Kim, “A review on security in smart home development,” *Int. J. Adv. Sci. Technol.*, vol. 15, pp. 13–22, 2010.
- [148] M. Xie, S. Han, B. Tian, and S. Parvin, “Anomaly detection in wireless sensor networks: A survey,” *J. Netw. Comput. Appl.*, vol. 34, no. 4, pp. 1302–1325, 2011.