Using Smart Homes to Detect and Analyze Health Events

Gina Sprint, Diane Cook, Roschelle Fritz, and Maureen Schmitter-Edgecombe
Washington State University, Pullman, WA 99164-2752

***Abstract*—** Instrumented smart homes offer an unprecedented opportunity to unobtrusively monitor human behavior in natural environments. Additionally, they can be used to determine whether relationships exist between behavior and health changes. Here we introduce an approach to behavior change detection (BCD) that can be used to identify behavior changes that accompany health events. BCD detects changes between time periods, determines significance of the detected changes, and analyzes the nature of the changes. In the case of smart homes, sensor data is collected and labeled using activity recognition and BCD is applied to analyze behavior changes by quantifying and analyzing changes in the activity timings and durations. We demonstrate our approach using three case studies for older adults living in smart homes who experienced major health events. Our evaluation indicates that behavior changes consistent with the medical literature do occur in these cases and that the changes can be automatically detected using BCD. The proposed smart home, activity recognition, and change detection algorithms are useful data mining techniques for understanding the behavioral effects of health conditions.

***Keywords*— pervasive computing, machine learning, time series analysis**

1. INTRODUCTION

In recent years, sensors have become ubiquitous in our everyday lives. Sensors are ambient in the environment, embedded in smartphones, and worn on the body. Data collected from sensors form a time series in which each sample of data is paired with an associated timestamp. This sensor-based time series data is valuable when detecting and analyzing changes associated with seasonal variations, new lifestyle choices, or new job situations. Analyzing sensor-based time series data can also be used to monitor changes in human behavior that are related to health events such as a fall, cancer treatment, or onset of a chronic medical condition. Automatically tracking behavior changes from sensor data can help with understanding the behavior impact of these health events. Similarly, detecting these changes can alert individuals and their caregivers about potential health concerns.

In this paper, we introduce a method to analyze the behavioral impact of health events using smart home sensor data called Behavior Change Detection, or BCD. Smart home sensor systems provide the capability to automatically collect information about a resident's everyday behavior without imposing any restrictions on their routines. We collect data from ambient sensors placed in smart home environments and label the data with the corresponding activities using automated activity recognition. To track changes in routine behavior, we quantitatively compare two or more time periods, or windows, of activity-labeled data. If the two time windows contain significantly different activity information then this may indicate a significant behavior change. In addition, we employ a virtual classifier to provide an explanation of the detected change.

To evaluate BCD, we analyze smart home data collected for multiple years in the homes of older adults. Health events are identified for three of the smart home residents based on medical records review and monthly interviews with the study participants. Data surrounding the health event is compared with baseline normal data to determine if a significant behavior change has occurred and describe the nature of the change.

The corresponding behavior change is then analyzed by a clinician to validate the behavior change and explain the relationship between the health event and corresponding behavior change.

Clinical studies support a relationship between daily behavior and cognitive and physical health [1]. Most of the prior work in this area utilized wearable data to correlate home-based movement with health measures [2], although smart home data has been used to analyze mobility and time out of the home with respect to cognitive and physical health [3], [4]. Our own earlier work showed that smart home data can be analyzed over time to predict performance on cognitive health assessment tests [5]. We hypothesize that the relationship between sensed behavior and health events can also be observed and analyzed using smart home data, which has not yet been examined. Results from the case studies presented in this paper indicate that smart home and machine learning technologies can be used to understand the behavioral impacts of health events and to provide information to individuals that indicate possible health concerns.

2. Collecting and Labeling Smart Home Data

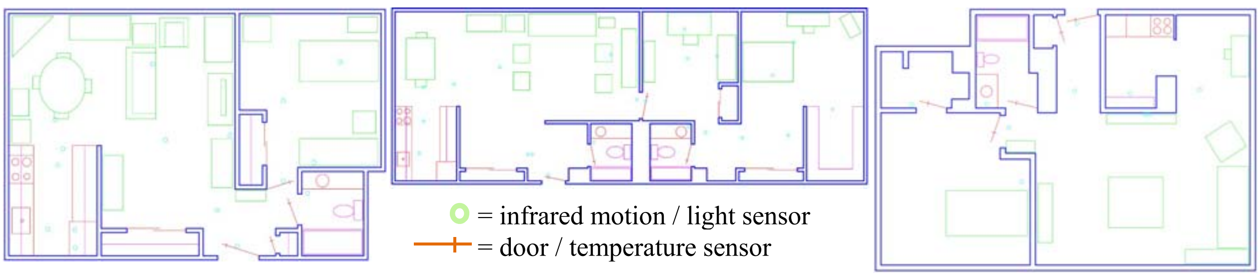

2014-06-15 03:38:28.094897 M009 ON		2014-06-15 03:38:28.094897 BedroomMotion ON Sleep
2014-06-15 03:38:29.213955 M009 OFF		2014-06-15 03:38:29.213955 BedroomMotion OFF Sleep
2014-06-15 03:38:17.814393 M015 ON	→	2014-06-15 03:38:17.814393 BathroomMotion ON Bed-Toilet
2014-06-15 03:38:58.584179 M015 OFF		2014-06-15 03:38:22.584179 BathroomMotion OFF Bed-Toilet
2014-06-15 03:39:17.814393 M009 ON		2014-06-15 03:39:17.814393 BedroomMotion ON Sleep

Figure 1. Smart home floorplan and sensor layout for three testbeds: SH1 (left, with 3 door/temperature sensors and 10 motion/light sensors), SH2 (center, with 5 door/temperature sensors and 23 motion/light sensors), and SH3 (right, with 3 door/temperature sensors and 10 motion/light sensors). Sample raw sensor data is converted to use generalized sensor identifiers and automatically labeled by CASAS-AR activity recognition with corresponding activity labels (bottom).

We collect data in everyday home environments using the CASAS "smart home in a box" [6]. The three homes that we include in this study are single-resident apartments, each with at least one bedroom, a kitchen, and a dining area. The apartment floorplans, sensor positions, and sample labeled sensor data are shown in Figure 1. These homes are equipped with combination motion/light sensors on the ceilings and door/temperature sensors on cabinets and doors. The sensors continuously and unobtrusively monitor daily activities of the residents by sending text message-type updates, or sensor events, whenever they sense a state change (i.e., from "door closed" to "door open" or from "no motion" to "motion"). The CASAS middleware collects these sensor events and stores them in a relational database.

Once the sensor data is collected we label each sensor event with the corresponding activity using the CASAS-AR activity recognition algorithm [7]. Let $A=\{a_1,a_2,..,a_T\}$ be the set of all activities. Given features $x \in \Re^d$ extracted from a sequence of sensor events ending at time *t*, the challenge of activity recognition is to map *x* onto a value $a \in A$ indicating the activity that occurred at time *t*. These labels provide a vocabulary for expressing and analyzing the sensed behavioral patterns. Activity recognition algorithms have been designed for wearable, phone, home, video, and other sensors using machine learning techniques that range from naïve Bayes classifiers and decision trees to more complex models including Gaussian mixture models and conditional random fields [8], [9].

AR is particularly well suited for this type of analysis because it does not require that the sensor data be pre-segmented into distinct activity sequences. Instead, it labels sensor events with activity labels in real time as the events occur. To do this, it moves a dynamic-size sliding window over the sensor events and extracts features *x* describing the current window of information. The features include the sensor event time of day, the size of the sliding window, the event count for each sensor within the window, time elapsed for each sensor since its most recent event, the most recent event location and sensor identifier, and the sensor generating the most events in the previous two windows.

Training data for CASAS-AR are provided by external annotators who look at one month of data and utilize both the house floorplan and resident information to generate corresponding ground truth activity labels [10]. In addition, sensor identifiers are replaced by more general location-based descriptors, as shown in Figure 1. Using this method, CASAS-AR learns an activity model based on training data from multiple smart home sites and can thus generalize for application to new smart homes with no training data. Although CASAS-AR has been tested with a number of classifiers including naïve Bayes, decision trees, hidden Markov models, and conditional random fields, the best performance was achieved using a decision tree. In this study we analyze the activities of Hygiene, Sleep, Bed-Toilet, Eat/Drink, Enter/Leave Home, Relax, and Work. For these activities in the three smart home testbeds we analyze in this paper, CASAS-AR achieved a recognition accuracy of 98% using 3-fold cross validation.

3. Detecting and Analyzing Behavior Change

We are interested in analyzing the behavioral impact of health events. More specifically, we want to determine if a significant change in behavior has occurred at the time of the health event and to analyze the nature of the behavior change. To do this, we introduce methods to quantify the amount of change in activity patterns between two windows of time series activity data that were sampled by smart home sensors and labeled by CASAS-AR. Let *X* denote a sample of time series data where each day's data are expressed by extracted activity features, $X = \{x_1, x_2, ..\}$ and let *W* be a window of *n* days such that $W \subseteq X$. For this study, activity features consist of the amount of time spent on each activity for a particular day and the sensor density of each activity (measured as number of sensor events) for a particular day. We also collect the total amount of movement that occurs in the home for the day, expressed as the total distance travelled by the person in the home. These features were shown in earlier work to provide insight on behavior patterns that correlate with cognitive and physical health of smart home residents [10].

BCD compares two windows of data, W_i and W_j, within time series *X*. In this paper the windows are one week in length (n=7) and BCD compares a baseline window ($i=1$, the first week in our data subset representing normal behavior for the resident) with each subsequent window ($j=2,3,\ldots$). We utilize three change detection methods. Each of these methods provides a slightly different perspective on the data comparison. Additionally, the more methods that detect a significant change, the greater is the evidence for a behavior change.

Method 1: RuLSIF. This is a non-parametric approach that determines the amount of change between two time series samples by comparing the probability distributions of the two samples. Instead of estimating the probability distributions which is computationally costly, we directly estimate their ratio. Relative unconstrained Least-Squares Importance Fitting (RuLSIF) [11] represents one such approach that estimates the ratio using the Pearson divergence dissimilarity measure.

RuLSIF does not explicitly provide a method to determine a cutoff threshold for Pearson divergence values that are considered significant change scores. To address this issue, we introduce a change significance test based on intra-window variability and outlier detection. The proposed change significance test utilizes the existence of day-to-day variability in human behavior patterns [12]. For a change between two windows to be significant, the magnitude of change (inter-window change) should exceed the day-to-day variability within each window (intra-window change).

To compute significance of the change score *CS* between two windows, we first generate a list of all possible daily change scores, *DCS*, within each window (there are 2 × Combination(n,2) such scores). Next, boxplot-based outlier detection is applied to see if *CS* is an outlier when compared to the distribution of intra-window daily changes scores, *DCS*. Here an outlier can be defined as an observation which appears to be inconsistent with other observations. To determine this, the interquartile range (75^{th} percentile – 25^{th} percentile) is computed. *CS* values outside of the 75^{th} percentile + 1.5 × interquartile range are considered outliers and thus significant. Advantages of this proposed significance test are that it is non-parametric and that it can be computed based on any window size.

Method 2: sw-PCAR. Our Permutation-based Change Detection in Activity Routine (PCAR) approach [5] was originally designed to analyze changes in longitudinal smart home data. Here we adapt the original approach to handle smaller windows of activity-labeled data. The resulting small-window PCAR (sw-PCAR) algorithm breaks each day within the window into non-overlapping hour-long time intervals. Each time interval has a corresponding probability distribution over the activities that occur at that time. For sw-PCAR, the days within two windows W_i and W_j are averaged to yield aggregate windows $\widehat{W_\iota}$ and $\widehat{W_J}$. Next, we compute a change score *CS* using the symmetric Kullback-Leibler (KL) divergence distance between the activity probability distributions in $\widehat{W_\iota}$ and $\widehat{W_J}$. Finally, the significance of the distance value *CS* is computed by concatenating data from windows $\widehat{W_\iota}$ and $\widehat{W_J}$ into one window *W*. All of the time intervals within *W* are randomly shuffled then split into two new sub-windows and the KL distance is computed for this permuted window pair. This shuffling procedure is repeated *N* times to produce a *N*-length vector *V* of KL distances. If *N* is large enough, the corresponding set of KL distances forms an empirical distribution of the possible permutations of activity data for the two windows. sw-PCAR computes change significance by comparing *CS* to the permutation vector *V* using boxplot-based outlier detection as we did with RuLSIF. If *CS* is identified as an outlier of *V* then the change score is reported as significant.

Method 3: Virtual Classifier. Our final method utilizes a binary classifier to detect and explain behavior change. This type of virtual classifier for change analysis was first proposed by Hido et al. [13]. For the VC approach, feature vectors from window W_i are labeled with a positive class and feature vectors from window W_j are labeled with a negative class. VC trains a decision tree to learn a boundary between the virtual positive and negative classes. The resulting average prediction accuracy based on k-fold cross validation is represented as p_{VC}. If a significant change exists between W_i and W_j then the average classification accuracy p_{VC} of the learner should be higher than the accuracy expected from random noise which is $p_{rand}=0.5$, the binomial maximum likelihood of two equal-length windows.

To determine the significance of the change between two windows, the inverse survival function of a binomial distribution is used to determine a critical value, $p_{critical}$, at which n Bernoulli trials are expected to exceed p_{rand} at α=0.05 significance. If $p_{VC} > p_{critical}$ then a significant change exists between windows W_i and W_j. If a change significance test concludes that the change score is significant, then for our health event study we would also like to explain the source of change. Typically, this requires computing features that summarize the data and provide a meaningful context for change and applying change tests for those specific features. One of the advantages of the VC, however, is that by utilizing a decision tree learner its output includes an explanation of the source of change without reliance on statistical tests. Upon detecting a significant change, the decision tree is retrained on the entire dataset and inspected to reveal the features that are the most valuable in discriminating between the two windows of data.

4. ANALYZING BEHAVIORAL IMPACT OF HEALTH EVENTS

We collected data in smart homes with older adult residents for multiple years. For each study participant we also recorded health events with their date and event type, based on medical records and monthly interviews with the participants. Here we describe three of these health events and utilize these case studies to illustrate the use of BCD.

4.1. Case 1: Radiation treatment

Case 1 focuses on an 86 year old female resident living in a smart home testbed that we refer to as SH1 (see Figure 1 for the smart home floorplans). Three months into the data collection, the participant was diagnosed with lung cancer and started radiation treatment during week W_{10}. We hypothesize that radiation treatment will have an observable and quantifiable impact on her behavior. To validate this hypothesis, we use BCD to compare one-week baseline of smart home activity data (W_1) with two other weeks. The first comparison is with another pre-event week, namely the week immediately following the baseline (W_2). The second comparison is with the first full week during which the individual underwent radiation treatment (W_{11}).

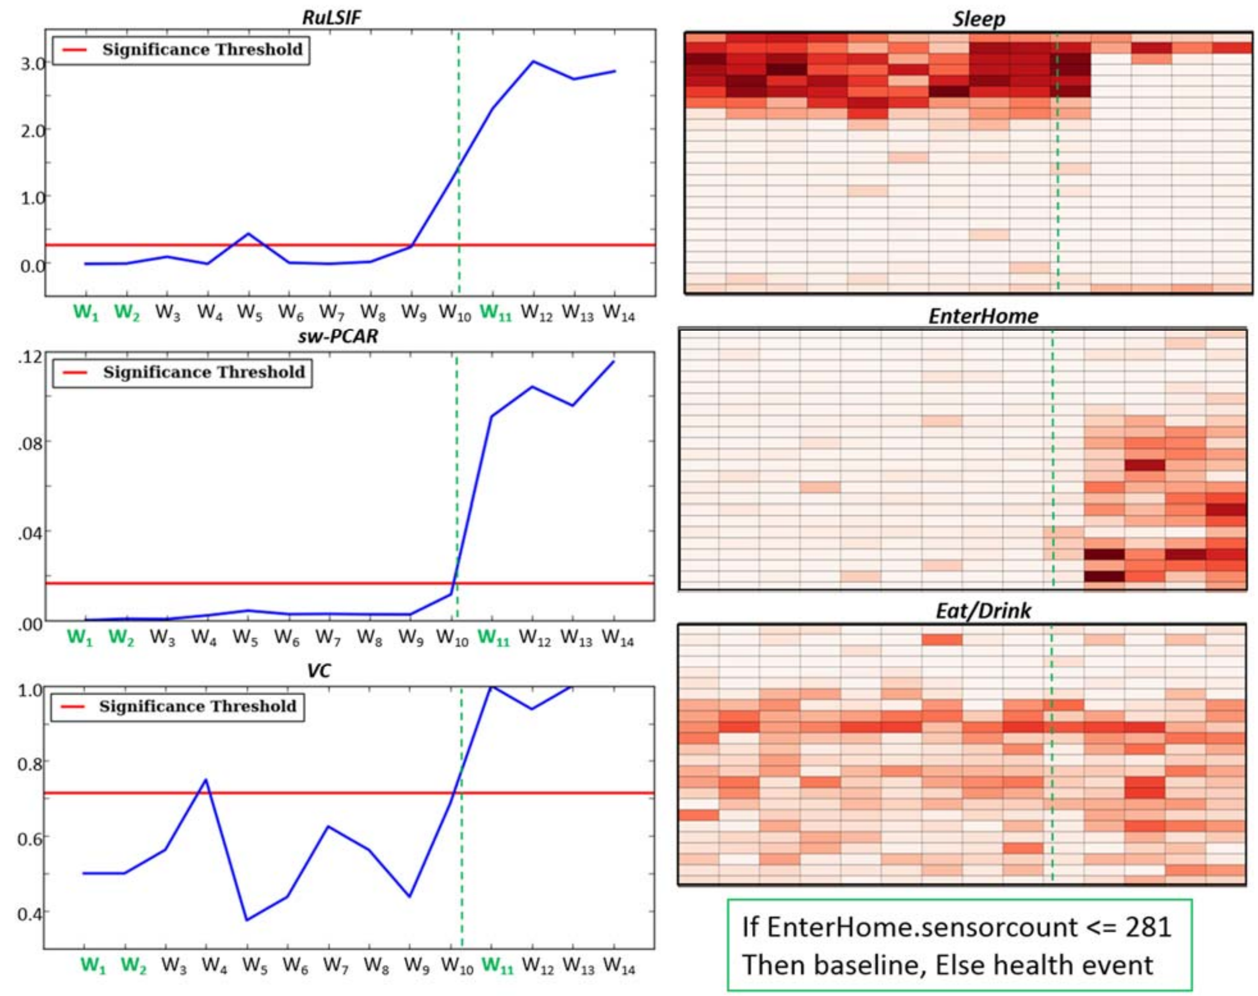

Figure 2. Results of SH1 health event analysis. Overall change scores are plotted (left) using RuLSIF, sw-PCAR, and VC comparing each week with the baseline week, W_1. Values above the red line show significant changes. Density maps for selected activities Sleep, EnterHome, and Eat/Drink are plotted (top right) for the same time period. Darker colors in the density maps indicate more time spent on the activity during that hour of the day. The top-level rule generated using VC is highlighted (bottom right) indicating the activity feature that best discriminates the baseline week from the health event week. In each plot the green dashed line indicates the occurrence of the health event.

Figure 2 illustrates results from applying each change detection method. Figure 2 also shows associated activity density maps for SH1. Density maps have been used in prior work to visualize levels of movement in the home [4]. Our activity density map is a heat map that visualizes the amount of time spent on a particular activity as a function of a 24-hour clock (y axis), aggregated over one week (x axis). The darker the color, the more time was spent on the activity during that particular hour of the day in the corresponding week.

As the density maps show, the participant's level of sleep decreased once treatment started and the number of times she left the home / returned home increased. Possible explanations for this are increased trips out of the home for appointments or visits from family and caregivers. Another impact of the treatment is the increased number of trips this participant made to the kitchen to eat or drink. These more frequent kitchen trips are consistent with the observation that radiation treatment increases the feeling of thirst, resulting in a patient drinking more liquids throughout the day [14].

Table 1. Change scores for smart home residents SH1, SH2, and SH3. Scores are computed between two normal activity weeks (W_1 and W_2) and between a normal activity week and a week during the health event (W_1 and W_{11} for SH_1 and SH2, W_1 and W_8 for SH3). For RuLSIF and sw-PCAR, larger values indicate greater change and values close to 0 indicate no change. In the case of VC, values close to 0.5 indicate no change and values close to 1.0 indicate large change. Significant results are indicated with an asterisk (*).

	Method	***W_1/W_2 (baseline)***	***W_1/W_{event} (health event)***
SH1	RuLSIF	-0.017	2.298*
	sw-PCAR	0.001	0.091*
	VC	0.500	1.000*
SH2	RuLSIF	0.010	3.315*
	sw-PCAR	0.004	0.042*
	VC	0.438	1.000*
SH3	RuLSIF	0.000	0.000
	sw-PCAR	0.000	0.001
	VC	0.500	0.750*

The change scores using the three BCD techniques described in this paper are summarized in Table 1. For participant SH1, the behavior changes during radiation treatment are evident for each of the change detection methods and the results are significant. The nature of the greatest change is highlighted by the decision tree that VC generates. As shown in Figure 2, the top-level feature is the number of sensor events that are related to an "EnterHome" activity. The number of times the participant (or a visitor) enters the home is larger during radiation treatment, with a great enough increase for this event to discriminate between baseline behavior and health event behavior.

4.2. Case 2: Insomnia

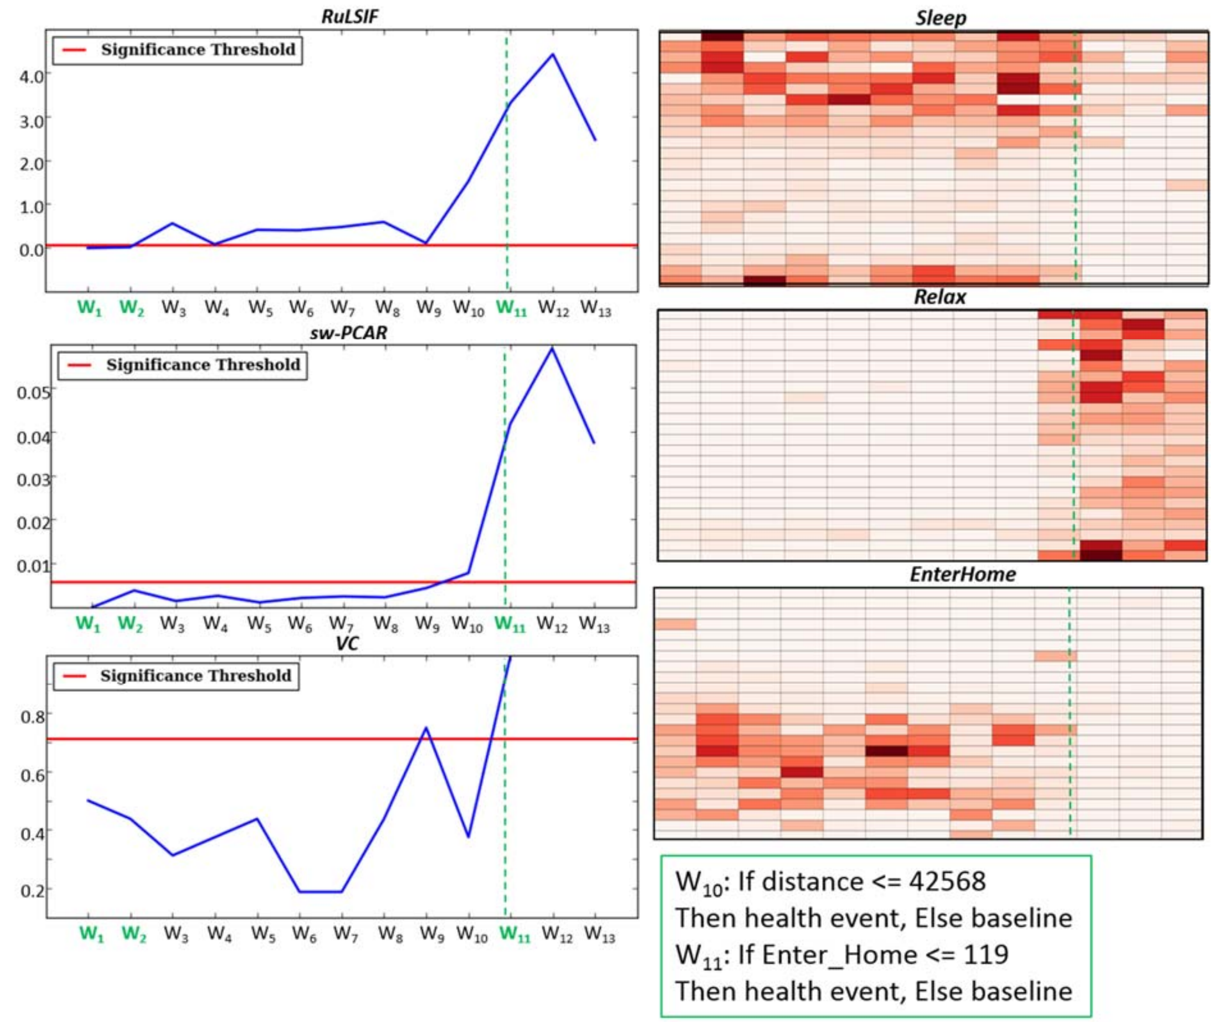

Figure 3. Results of SH2 health event analysis. Overall change scores are plotted (left) using RuLSIF, sw-PCAR, and VC for baseline week W_1 and health event week W_{11}. Values above the red line show significant changes. Density maps for selected activities Sleep, Relax, and EnterHome are plotted (top right) for the same time period. The VC-generated rule is also shown (bottom right). In each plot the green dashed line indicates the occurrence of the health event.

Case 2 is a 91 year old female smart home resident whom we refer to as SH2. During the time that data was collected in this participant's home, she was diagnosed with insomnia. To measure the impact of this health event on her sleep and on other routine activities, we use BCD to compare two weeks of normal behavior (weeks W_1 and W_2) and one week of baseline behavior with a week surrounding the insomnia diagnosis (weeks W_1 and W_{11}). The change scores are summarized in Table 1 and indicate that significant changes in overall routine are detected by all three methods.

In Figure 3 we see that changes occur not only during week W_{11} but in the days leading up to the health event and persisting to days and weeks following the insomnia diagnosis. We can also observe in the density maps that the amount of sleep does decrease during this period. The change in behavior also impacts relaxation, which is time spent in a favorite chair or couch with little movement and possibly napping. These relaxation periods occur during normal sleep hours but also throughout the day. In addition, the number of trips outside the home decreases during this time. The virtual classifier actually finds the corresponding decrease in EnterHome events to be the main discriminating feature between baseline and health event weeks. On the other hand, if we look slightly earlier at week W_{10}, VC again detects a significant change from the baseline week and the main discriminating feature is the total movement in the home throughout the day (measured as distance traveled in the home). This is a factor that could be considered when examining possible reasons for insomnia or impact of a decrease in sleep time.

4.3. Case 3: Fall

The last case, called SH3, is an 80 year old female living in a smart home testbed. During the time that we were collecting sensor data in this home, the participant fell in her home. She described that her right leg hurt for several days after that and "slowed her down". To analyze the impact of this health event we compared data collected at baseline (W_1) with the following week which also contained normal activity and no health events (W_2). We also compared W_1 with the week containing the health event (W_8).

As the results in Table 1 indicate, this health event has a subtler impact on behaviors, at least those that can be detected by ambient smart home sensors. RuLSIF and sw-PCAR detect almost no change between weeks W_1 and W_2 or between weeks W_1 and W_8. The virtual classifier is the only method that finds the change during the health event week. As the VC-generated rule indicates, the difference is primarily detected based on the total distance that the individual traveled throughout the home on a daily basis. The decrease in movement is consistent with the observation that the hurt leg caused the resident to slow down. As the density plots indicate, there appears to be less impact on other routine activities such as sleep and bed toilet transitions. There is an apparent slight decrease in trips out of the home but this is not large enough to be detected by the change detection methods.

5. Conclusions

In this paper we introduce BCD, an approach to behavior change detection. We describe how BCD can be used to quantify and explain changes that are detected in daily activity data. In particular, BCD can detect changes in smart home-detected behavior data that occur as a result of health events. From the three case studies that we analyzed in this paper we see that the ability to detect behavioral impact of health events depends on the nature of the health event itself. Some events impact multiple activities including sleep, eating, and trips out of the home. In contrast, other events have more localized impact. The ability to detect the actual health event occurrence (e.g., fall) and its impact may require additional, more sensitive sensors to be placed in the home or on the body. A systematic comparison of different BCD window sizes may also provide insights on the typical duration of behavior changes that may be associated with different types of health events. Future work also consists of analyzing all of the BCD-detected changes to determine the broader spectrum of events that elicit changes, such as failed sensors or visitors in the home.

The ability to detect behavior changes that are associated with health events is valuable for researchers who want to better understand the relationship between health and behavior. These insights may also help care providers respond to the needs of individuals who are experiencing changes in their health. An algorithm such as BCD can periodically look for changes in behavioral routine and alert the individual and their caregiver about these changes as they may indicate changes in cognitive or physical health. Because BCD can analyze any type of sensor data, our continued research will adapt these methods to analyze smart phone and wearable data, as well as data collected in smart homes.

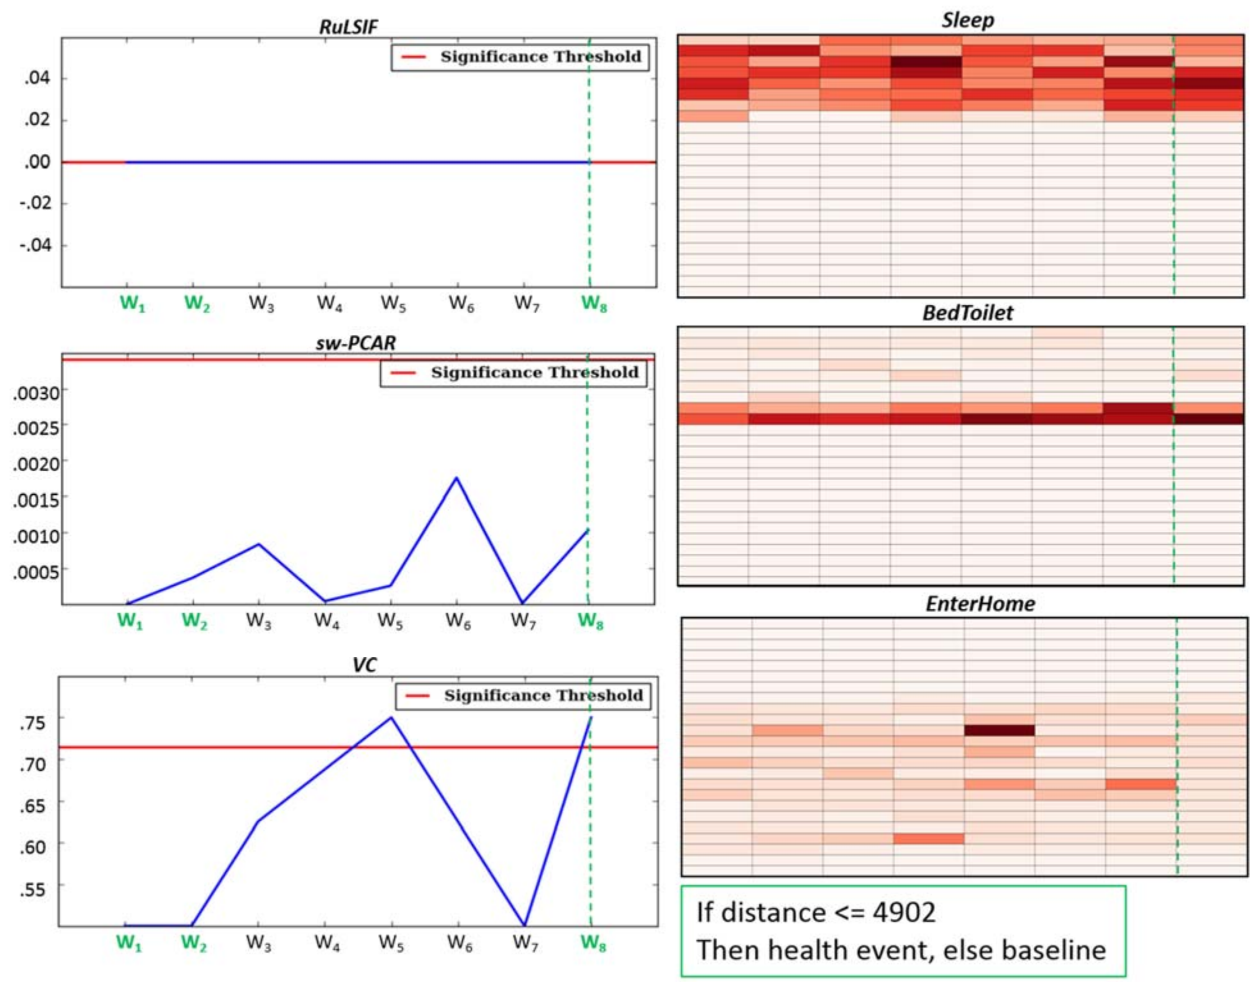

Figure 4. Results of SH3 health event analysis. Overall change scores are plotted (left) using RuLSIF, sw-PCAR, and VC for baseline week W_1 and health event week W_8. Values above the red line show significant changes. Density maps for selected activities Sleep, BedToilet, and EnterHome are plotted (top right) for the same time period. The VC-generated rule is also shown (bottom right). In each plot the green dashed line indicates the occurrence of the health event.

ACKNOWLEDGMENTS

This work is funded in part by National Science Foundation grant 0900781 and National Institutes of Health grant R01EB015853.

References

[1] I. M. et al. Lee, “Effect of physical inactivity on major non-communicable diseases worldwide: An analysis of burden of disease and life expectancy,” *Lancet*, vol. 380, pp. 219–229, 2012.

[2] 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.

[3] S. Robben, M. Pol, and B. Krose, “Longitudinal ambient sensor monitoring for functional health assessments,” in *ACM International Joint Conference on Pervasive and Ubiquitous Computing*, 2014, pp. 1209–1216.

[4] S. Wang, M. Skubic, and Y. Zhu, “Activity density map visualization and dissimilarity comparison for eldercare monitoring,” *IEEE Trans. Inf. Technol. Biomed.*, vol. 16, no. 4, pp. 607–614, 2012.

[5] P. Dawadi, D. Cook, and M. Schmitter-Edgecombe, “Modeling patterns of activities using activity curves,” *Pervasive Mob. Comput.*, 2015.

[6] 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.

[7] N. Krishnan and D. J. Cook, “Activity recognition on streaming sensor data,” *Pervasive Mob. Comput.*, vol. 10, pp. 138–154, 2014.

[8] 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.

[9] 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.

[10] P. Dawadi, D. J. Cook, and M. Schmitter-Edgecombe, “Automated clinical assessment from smart home-based behavior data,” *IEEE J. Biomed. Heal. Informatics*, 2016.

[11] S. Liu, M. Yamada, N. Collier, and M. Sugiyama, “Change-point detection in time-series data by relative density-ratio estimation,” *Neural Networks*, vol. 43, pp. 72–83, 2013.

[12] R. Refinetti, G. C. Lissen, and F. Halberg, “Procedures for numerical analysis of circadian rhythms,” *Biol. Rhythm Res.*, vol. 38, no. 4, pp. 275–325, 2007.

[13] S. Hido, T. Ide, H. Kashima, H. Kubo, and H. Matsuzawa, “Unsupervised change analysis using supervised learning,” in *Advances in Knowledge Discovery and Data Mining*, 2008.

[14] P. Beach, B. Siebeneck, N. Buderer, and T. Ferner, “Relationship between fatigue and nutritional status in patients receiving radiation therapy to treat lung cancer,” *Oncol. Nurs. Forum*, vol. 28, no. 6, pp. 1027–1031, 2001.