The 5th IEEE International Conference on E-Health and Bioengineering - EHB 2015
Grigore T. Popa University of Medicine and Pharmacy, Iaşi, Romania, November 19-21, 2015

Early illness detection in elderly using sensor networks: a review of the TigerPlace experience

Mihail Popescu, IEEE Senior Member
University Of Missouri, Health Management and Informatics Dept.
Columbia, Missouri, USA

Abstract- ***Many older adults in the US prefer to live independently for as long as they are able, despite the onset of conditions such as frailty and dementia. Elderly patients are particularly at-risk for late assessment of health changes due to factors such as their impression that such changes are simply a normal part of aging or their reluctance to admit to a problem. In-home sensors networks have emerged in the last ten years as a possible solution for early illness detection. Many projects have demonstrated the utility of in-home sensors for monitoring elderly but also have shown the necessity of developing new pattern recognition algorithms able to handle large amounts of diverse data. In TigerPlace, an aging in place facility from Columbia, MO, we created a unique living laboratory by deploying in-home sensors together with an electronic health record (EHR) system developed in-house that integrates clinical and sensor data. In-home monitoring devices such as infrared motion detectors, Kinect depth cameras, Doppler radars and bed sensors capture information related to the behavior of the residents from the monitored apartment and assist the clinical personnel in medical decision making. In this paper we present a review of our early illness detection (EID) and recognition (EIR) methodologies experimented in TigerPlace, together with results and lessons learned.***

Keywords: sensor networks, early illness detection, sensor sequence similarity, frequent activity patterns.

I. INTRODUCTION

Between 2015 and 2050, the USA will experience a considerable increase in population aged 65 and over, due largely to baby boomers who turned 65 in 2011 [1]. Factors such as fertility decline (fewer available children to live with), increased income and an individual centric modern culture make older adults chose to live independently despite advanced age and the onset of chronic conditions such as chronic heart failure or dementia [2]. Some chronic disease and their exacerbations are preventable or less costly if they are diagnosed in early stages by clinicians with help from family members [3]. According to a Center for Disease Control (CDC) overview of chronic disease spending [4], the cost of the five most prevalent chronic diseases in US (heart disease, cancer, diabetes, arthritis and obesity) was close to a trillion dollars. Consequently, early illness detection (EID) may lead to a better quality of life and a reduction in health care cost.

Sensor networks have emerged in the last decade as a possible solution to EID in older adults. Many academic projects such as CASAS (Washington State University), TigerPlace (University of Missouri) and ORCATECH (Oregon Health and Science University) have demonstrated the utility of in-home sensors for monitoring elderly but also have also shown the necessity of developing new pattern recognition algorithms able to handle large amounts of diverse data also known as big data.

In TigerPlace [5][6], an aging in place facility from Columbia, MO, our interdisciplinary research team has created a unique living laboratory by deploying in-home sensors together with an electronic health record (EHR) system developed in-house that integrates clinical and sensor data. In-home monitoring devices such as infrared motion detectors, Kinect depth cameras, Doppler radars and bed sensors capture information related to the behavior of the residents from the monitored apartment and assist the clinical personnel in medical decision making. We have shown that unobtrusive, continuous monitoring of individuals with in-home sensors provides useful embedded health assessment that can assist health care providers in detecting health decline and leading to earlier intervention [7,8].

A variety of methodologies for detecting activity and assess medication compliance have been reported in the literature [9-14]. In TigerPlace, we employ sensor network technology to provide early illness detection. We have installed sensor networks in the apartments of 50 residents, a system that has been active since fall of 2005. Using the deployed sensor network we experimented with various EID and EIR methodologies. Here, we distinguish between EID methodologies that signal a possible onset of a non-specific illness and EIR ones that try to provide more details and even identify the medical condition. The EID methodologies detect if a day is normal (no clinical significant event was detect) or abnormal (some clinical significant was detected). The EID algorithms can be supervised, such as k-nearest neighbor-(KNN) or unsupervised, such as one class classifier (OCC). EIR methodologies try to identify possible diseases associated to the captured behavior such as: loss of appetite and lethargy may be due to gastroenteritis, frequent nightly toilet visits might be due to a urinary tract infection (UTI) and lethargy and decreased walking speed may due to depression [14]. While a supervised approach to EIR is theoretically possible, it is most of the time impractical to find behavior patterns for all known diseases and even harder to use them across patients. Our EID method is based on linking the sensor data to the TigerPlace EHR data using a bioinformatics approach. In this paper we present seven EID/EIR methodologies: A) anomaly detection based on a single sensor, B) anomaly detection based on multiple sensors, EID using KNN, C) anomaly detection using sequence similarity, D) anomaly detection based on missing frequent normal patterns (MFP), E) anomaly detection based on frequent abnormal patterns, F) EID based on sensor similarity and EHR data, and G) blood pressure variation prediction.

The structure of this paper is as follows. In section II we briefly describe the sensor network architecture and the data it produces. In section III we describe two sensor representation methodologies used in our research: multidimensional time series used in methods 1, 2 and 7 and unidimensional discrete time series that is related to the methods 3, 4, 5 and 6 mentioned above. In section IV we present seven EID/EIR methods tested on TigerPlace data. Finally, in section V we provide conclusions and future directions.

II. System Architecture

We deployed our sensor networks in 50 TigerPlace apartments with the University of Missouri IRB approval. On average, we have two years of data for each resident. Figure 1 shows the schema of our monitoring system.

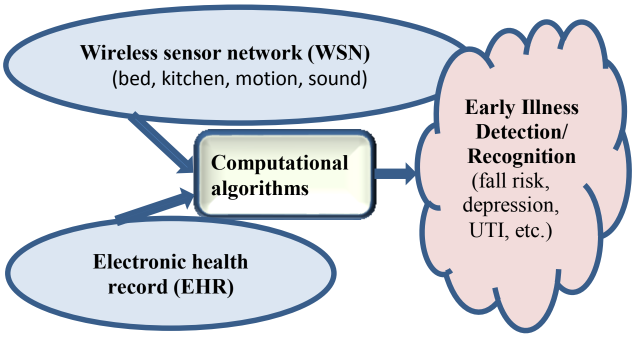

Fig. 1. TigerPlace monitoring system architecture.

The main components of the monitoring system are: a sensor network, a data logger, a reasoning system that analyses sensor patterns, an electronic health record (EHR) system, an alert manager to notify clinicians of potential problems, and a secure Web-based interface to display the data for the clinicians and researchers. The data logger date-time stamps the data, and sends them to a database on a secure server via a wired network connection. The passive infrared (PIR) motion sensors are placed in various places, such as the living room, kitchen, bedroom, bathroom, on the ceiling over the shower, in the laundry closet, in the refrigerator, and in kitchen cabinets and drawers. Bathroom activities are monitored by a motion sensor installed above the shower. The depth camera (Microsoft Kinect) is placed in the living room and measure gait parameters (speed and step length) and detects falls. The sleep patterns of each resident are captured by a bed sensor placed under the mattress. The bed sensor consists of four hydraulic strips [12] that captures quantitative heart rate, respiration and restlessness values.

The EHR system captures a variety of data about a TigerPlace resident such as demographics, International classification of disease version 9 (ICD9) diagnose codes, medications, emergency room visits, hospitalization records, nursing progress notes, nursing visit notes, assessment forms (SF-12, mood scale, fall assessment, mini mental state exam-MMSE), interim physician orders (IPO), reimbursement forms (HCFA-485) and activities of daily living (ADL). The nursing visit module records information about vital signs and certain assessment questions; health data are entered into the EHR system during the weekly wellness visit or during visits requested by the resident. The progress note module is used by TigerPlace staff to share information about a resident. The ADL module assists the nursing staff to record and assess daily activities of each resident. As part of the EIR framework, the EHR data is linked to the WSN database so that sensor data can be automatically annotated using health data.

III. Sensor Sequence Representation Methods

In TigerPlace we used two sensor sequence representations: a multidimensional time series made of aggregated hit values and a unidimensional discrete sensor sequence made of individual time-stamped hits.

A. *Multidimensional time series (MTS) representation*

The MTS representation was based on aggregating the hits of each sensor in a certain interval (15 min, hour or part of a day). The dimension of the time series was given by the number of sensors considered. Sometimes, we aggregated the firings of a set of sensors (e.g. all motion sensors in the apartment) to reduce the dimensionality of the resulting time series. An example of daily firing sequence for all PIR motion detectors from a TigerPlace apartment is shown in Figure 2.

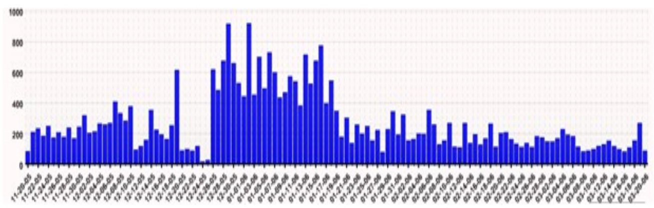

Fig. 2. A typical daily motion sensor firing sequence: each bar represents the daily sum of motion firings in the entire apartment for a given resident.

B. *Unidimensional time series (UTS)*

Our UTS representation of sensor sequences has a bioinformatics inspiration. Formally, a discrete sensor sequence can be defined as $T_1=\{(C_1\ t_1), (C_2\ t_2), \ldots, (C_m\ t_m)\}$ where C_i are a set a symbols from an alphabet Σ associated to the deployed sensors and t_i is the time when firing C_i was recorded. Several examples of representations for symbolic sequences can be seen in Figure 3.

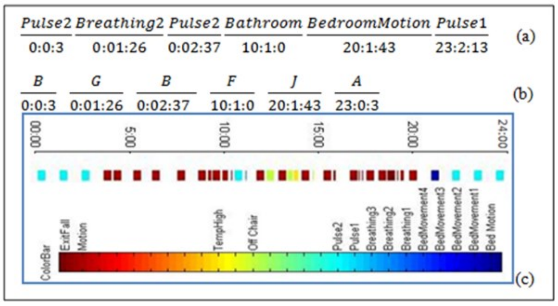

Fig. 3. a) An example of sensor sequence with 6 firings; b) same sequence represented using a user defined alphabet; c) a more complicated sequence represented using a user defined color code.

For sensors that have native discrete firings, such as the PIR motion sensors (see *Bathroom* and *BedroomMotion* firings in Figure 3.a), this representation is natural. However, for sensors that produce continuous values (such as Pulse and Breathing) we need to provide a discretization scheme such as: *Breathing1* = Breathing < 10 breaths/min, *Breathing2*= Breathing < 30 breaths/min and *Breathing3*= Breathing > 30. Similarly Pulse and Bed Restlessness were discretized using three and four intervals as [0, 40, 100, …] beats/min and [0,4,7,15,…] s, respectively. For example, *Pulse2* fires when Pulse is between 40 and 100 beats/min and *BedMotion3* fires when the resident moves in his bed between 7 and 15 s. The advantage of this discrete representation is that it avoids defining a multi-dimensional time series distance which might be a complicate problem mainly when the number of sensors is large. Moreover, it avoids aggregating hits in certain time intervals (hourly for example) which can improve prediction accuracy.

Note that in Figure 3.b the sensor representation looks like an amino acid sequence with two differences: it has an associated time (symbols are not equidistant), and it may have more than 20 symbols. In fact, if the size of the alphabet Σ is less or equal to 20, one can directly use bioinformatics tools such as sequence similarity (BLAST, Smith Waterman) and motif finders (MEME), by ignoring the temporal dimension.

For considering the temporal dimension, in [15, 16] we proposed a temporal Smith Waterman (TSW) algorithm for computing the similarity between two discrete sensor sequences.

IV. EID/EIR METHODS IN TIGERPLACE

A. EID: Health alerts based on single sensor values

Using the MTS representation described in section III.A we computed the distribution of the normal sensor values for each dimension of the time series [8, 17]. For sensor i we assume that the values in the near past (two weeks) follow a 1-D Gaussian distribution with mean μ_i and standard deviation σ_i. If a recorded value for sensor i is a number n of standard deviation σ_i away from the mean, we send an alarm to the clinical personnel. For example, in TigerPlace, we used n=2 for respiration, n=3 for restlessness and n=4 for pulse. The alerts were evaluated using feedback from the clinicians [7, 8]. The clinician receives an email with two links: one to the full data display and another to the evaluation system where the alert is rated on a scale from 1 (not-useful) to 5 (useful). In TigerPlace experiments [17] this method had an area under the receiver operating curve (AROC) of about 0.4. This methodology is in fact an one-class classifier (a.k.a. anomaly detection) approach since it uses only sensor values from one class (normal days) to compute the Gaussian distribution. Aside of the poor performance, this alert methodology tends to send multiple messages for the same event, leading to a clinician information overload.

B. EID: Health alerts based on multiple sensor values

More complicated one-class classifiers (OCC) such as nearest neighbor (OCCNN), support vector machine (OCSVM) or Gaussian mixture (OCGM) [18,19] can be used if more sensor dimensions are employed. For example, in [19] we used 4 sensor values, namely, low pulse (*Pulse1*), low breathing (*Breathing1*), bed restlessness (*BedRestlessness1*), and overall motion with OCCNN and OCSVM for EID. The feature vector had 8 dimensions: features 1-4 were the sum of the 4 sensor firings for the night hours (7pm-7am) and features 5-8 were the sum for the day hours (7am-7pm). On a pilot dataset, OCCSVM proved to be better than OCCNN ($AROC_{OCCSVM}$ was about 0.65, $AROC_{OCCNN}$ about 0.55). Moreover, both multidimensional OCCs were better than the 1-D OCC described in section IV.A (AROC about 0.4).

C. EID: Health alerts using sensor sequence similarity

An alternative anomaly detection strategy is to exploit the similarity between sensor sequences when using UTS representation [20]. The main idea of this approach is that sensor sequences cased by abnormal behavior are probably very dissimilar from the one recorded in recent resident past. The key concept of this approach is a similarity measure between sensor sequences like the TSW [15, 16] mentioned in section III. Given n (past) normal sensor sequences $\{S_i\}_{i=1,n}$ we compute the pair-wise similarities between them, $\{s_{ij}\}_{i,j=1,n}$ using TSW [20]. We, then, calculate the distribution of the $\{s_{ij}\}$ similarities of these "normal" days assuming they follow a Gamma distribution. Example of daily sensor sequence similarity distributions for three TigerPlace residents can be seen in Figure 4.

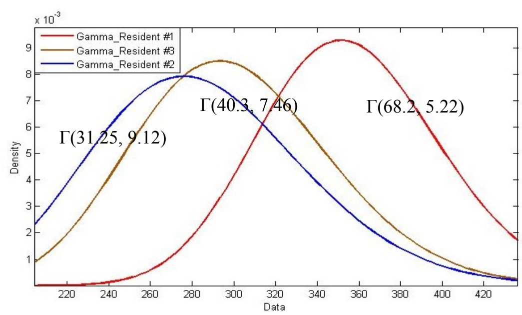

Fig. 4. The Gamma distribution of the daily similarity for three TigerPlace residents [20].

The distributions look somewhat Gaussian, although better results were obtained using a Gamma distribution instead. Assume that we found that $\{s_{ij}\}$ follow a Gamma distribution with parameters a and b, i.e. $\Gamma(s_{ij},a,b)$. To classify an unknown sequence S_x, we start by computing its similarities with all normal sequences $S_1,\ldots S_n$, obtaining similarities $\{s_{ix}\}_{i=1,n}$. Then, we find their maximum $s_x=\max\{s_{ix}\}$. The confidence that S_x is abnormal, $C(S_x)$ is 0 if $s_x > a/b$ and equal to 1-P where $P=\Gamma(s_x,a,b)$.

This approach resulted in an AROC performance of about 0.75 on our pilot datasets. Anomaly detection based on sensor similarity may lead to better results than the previous OCC approach due to the sensor representation method, UTS vs MTS, as sensor aggregation in MTS may lead to information loss. In the same time, the choice of distance in a MTS space is still a difficult problem due to sensor value heterogeneity that requires a trial-and-error approach to finding the best one.

D. EID: anomaly detection based on missing frequent normal patterns (MFP),

The weakness of the above approach is that is based on day long (24 h) sequences, S_i, which are not granular enough to describe the activities of daily living (ADL) of the TigerPlace residents. The question is: can we find those sub-sequences (denote here as "motifs") that correspond to frequently performed activities, also denoted as frequent paterns (FP)? While activity recognition based on some classifier such as SVM, Hidden Markov Model would result in better activity recognition performance, this approach is not suitable for monitoring systems, where it is hard to collect the data and train a system for each monitored person. Instead, we used a similar approach to the bioinformatics "frequent motif discovery" used in algorithms such as MEME [21]. In the first approximation, we neglected the time between firings, so our sequences were identical to amino acid ones. In Figure 5.a [22] we show an example of a daily sequence in Fasta format, used for MEME [21] input. In Figures 5.b and 5.c we show 2 bathroom FAPs computed by MEME represented using a sequence logo approach: the size of each letter (sensor) is proportional to the probability of that sensor firing in the given location. We see that in some days (5.c), the resident showers less than in others (5.a): "S" is smaller in 5.c than in 5.a.

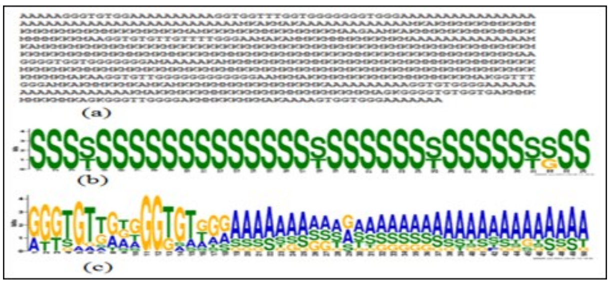

Fig. 5. (a) Example of input to the MEME tool: a day-long sensor sequence in fasta format where each letter denotes a sensor; (b) example of MEME output: shower activity motif. S=shower sensor, T=bathroom sensor, G=cabinet sensor; (c) example of MEME output: night bathroom visit. A=living room motion sensor. [22]

Other methods, such as frequent T-patterns [23], can be used to find FPs in temporal sequences. Activity recognition experiments performed on bathroom motifs extracted from 3 TigerPlace residents achieved an AROC of about 0.8, which is a 5% improvement over the approach based on a day-long sequence.

To perform EID using a sensor network, we need to identify FPs in recent past data (say last two weeks). Note that we don't need to know the identity of the FPs (i.e. what activity they represent), just that they repeat as part of the resident normal daily/weekly routine. The missing frequent pattern (MFP) framework [22] assumes that if a given day is missing some number of FPs, then the resident might be facing a health problem. MFP method is another take on anomaly detection and it is unsupervised since we don't need to know which days are abnormal or the identity of the frequent patterns.

E. EID: anomaly detection based on abnormal frequent activity patterns (AFAP),

The AFAP framework is similar to the MFP one, but instead of using frequent normal patterns, it is based on abnormal frequent abnormal patterns, that is FP extracted from abnormal days. AFAP and MFP can be run together synergistically: first MFP is used to find abnormal days. These days can in turn used to extract FP related to abnormal behavior. After a significant number of abnormal FPs are accumulated, they can be further used to detect abnormal days. On a retrospective study on three TigerPlace residents [23] using the AFAP framework, we obtained an AROC of about 0.7. A day was declared abnormal if three abnormal patterns would be found in that day. The patterns were identified using a TSW similarity with a threshold of 0.9.

F. EIR: based on sensor similarity and EHR data [16]

Sequence similarity based inference methods have been extensively used in bioinformatics. For example, a "guilt by association" approach (also known as "annotation") is used to find molecular functions for unknown gene products by comparing their amino acid sequence to sequences of known proteins. In our case, given a database of sensor firings and related health data (from the integrated EHR and WSN, see Figure 1) we can infer current health events that have been associated in the past with sequences similar to the one in question [15]. Formally, given M pairs of sensor firing sequences (from WSN) and their associated conditions (from EHR), $\{S_i, C_i\}_{i=1,M}$, we can infer the conditions associated to

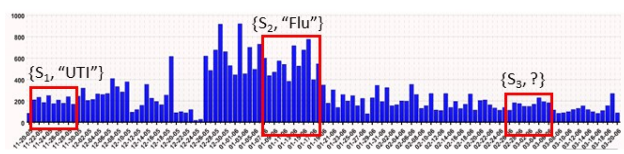

Fig. 6. EIR using sensor sequence similarity: since S3 is more similar to S_1 than S_2, than S_3 is most likely "UTI".

an unknown sensor firing sequence S_u by computing the similarity $s_i=\mathrm{sim}(S_u, S_i)$ using for example the TSW similarity measure mentioned in section III. That is, we associate the unknown sequence S_u with medical conditions for which the related similarity s_i was greater than some threshold $\varepsilon \in (0,1)$. We illustrate the idea in Figure 6, where the cause of the behavior captured by S_3 might be urinary tract infection (UTI), since S_3 is most similar to S_1 and not S_2.

Since this framework uses EHR data that is usually acquired when a nursing visit (hence an abnormal behavior) happens, it is complementary to the previously described methods. That is, knowing that the day is abnormal, what else can we say about it?

In [15] we showed some results obtained on a pilot dataset from three TigerPlace residents. The characteristics of the dataset are shown in Table I.

Table I. Characteristics of the TigerPlace pilot dataset

Resident #	# of sensor days	# of comments	# abnormal days
#1	440	83	81
#2	745	44	35
#3	500	499	335

We extracted the medical terms from the EHR nursing comments using Unified Medical Language System (UMLS) MetaMap (http://www.nlm.nih.gov/research/umls/) and mapped them to UMLS concepts. Each day was represented using a time series or sensor hits (UTS). The similarity between days was computed using TSW. For the above dataset we obtained an annotation average precision of 0.64 and a recall of 0.37.

Using a day-long sequence to capture the link between the sensor data and nursing comments is too coarse: it may happen that the behavior related to the illness present in the nursing notes was not present the entire day. We illustrate the problem in Figure 7.

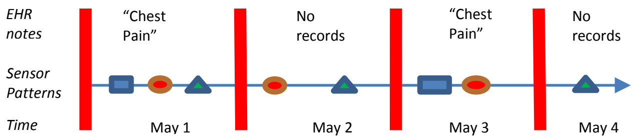

Fig. 7. Uncertainty of behavior location: it is hard to tell which pattern caused the symptoms captured in the EHR notes "chest pain". It is probably the "rectangle" pattern.

In Figure 7, we show multiple frequent patterns associated with various EHR concepts. Whereas we have the same EHR concept on May 1 and May 2, it would be incorrect to associate "Chest pain" with all the patterns from those days (rectangle, oval and triangle). In our example, the most probable FP associated to "Chest pain" is the rectangle FP.

In [25] we showed a possible solution to this problem based on multiple instance learning [MIL]. MIL is a supervised learning method in which individual labels for each training example are either hard to assign, such as labeling people in an image, or not available, such as in which hour of the night the resident didn't feel well. Instead, it is much easier to obtain labels for sets of objects (in our case a day) called bags, and then labeling the entire bag with a given label. In [25] we only used normal/abnormal labels, but MIL can be extended to multiple labels (classes).

G. EIR: blood pressure variation prediction.

In elderly, some medical condition such as hypertension or diabetes have a great influence on their behavior. In such cases, it is possible to correlate the sensor data to the medical condition. For example, in [24] we used bed and motion sensor data to predict the pulse pressure (the difference between systolic and diastolic pressures) trend in elderly residents with hypertension. We conducted a retrospective pilot study on two residents of the TigerPlace aging in place facility, with age over 70 and blood pressure measured between 100 and 300 times during a two year period. The blood pressure values were manually extracted from the nursing visit reports present in our EHR. A robust regression model was used to compute the pulse pressure (see Figure 8) based on the total number of motion and bed sensors hits during the night and day time, respectively (4 features). The pilot study suggested that pulse pressure trends can be reasonably well estimated (average relative error of less than 10%) using a bed and PIR motion sensors.

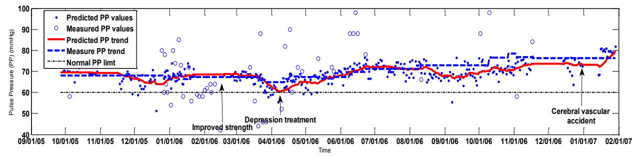

Fig. 8. Comparison between the computed (red) and measured (blue) pulse pressure (systolic-diastolic) for a male resident [24].

V. Conclusions

In this article we reviewed several early illness recognition/detection algorithms experimented in TigerPlace, our living laboratory from Columbia, MO, USA. Each methodology has its own positive and negative aspects. To get an idea of the performance of each method, we provided an approximate value of the AROC obtained on some TigerPlace dataset. Since the datasets were not identical for all methods, the value of the AROC has to be taken with a grain of salt.

Most of the above methods complement each other, hence the best EID strategy would be to use more than one at the same time together with some fusion methodology. For example, one can use multiple anomaly detection methodologies to find if a day is normal/abnormal and then use a voting scheme to decide the final label.

Treating sensor data as sequences of symbols shows promise for EID/EIR and, more general, for activity recognition. Unlike bioinformatics sequences, the sensor ones have a temporal dimension. For example, when we used MEME, we ignored time (see figure 5). Consequently, new methods have to de developed for representation and detection of temporal sequence motifs. As bioinformatics motifs are usually conserved across individuals and even species, the question arises if it is possible to identify similar (i.e. "conserved" as they are known in bioinformatics) activity motifs across individuals. This will greatly enhance the possibility of algorithm training and understand disease behavior in elderly.

ACKNOWLEDGMENT

Part of this work has been funded by a NSF SHWB grant, award #: IIS-1115956.

REFERENCES

[1] Ortman JM, Velkoff VA, Hogan H, "An Aging Nation: The Older Population in the United States. Population Estimates and Projections", http://www.census.gov, May 2014, P25-1140.

[2] Esbaugh EM, "Perceptions of Living Alone Among Older Adult Women", Journal of Community Health Nursing, 25:125–137, 2008.

[3] Boockvar KS, Lachs MS. Predictive value of nonspecific symptoms for acute illness in nursing home residents. Journal of the American Geriatrics Society. 2003; 51(8):1111–5.

[4] CDC, http:// http://www.cdc.gov/chronicdisease/overview/, accessed July 8 2015.

[5] MJ Rantz, K Dorman-Marek, M Aud, HW Tyrer, M Skubic, G Demiris, et al., "A technology and nursing collaboration to help older adults age in place", Nursing Outlook, 2005, pp.40-45.

[6] M Skubic, G Alexander, M Popescu, M Rantz, J Keller, "A Smart Home Application to Eldercare: Current Status and Lessons Learned", Journal of Technology and Health Care, 2009, pp.183-201.

[7] Rantz MJ, Skubic M, Koopman RJ, Alexander GL, Phillips L, Musterman K, Back J, Aud MA, Galambos C, Guevara RD, Miller SJ.(2012). Automated technology to speed recognition of signs of illness in older adults. Journal of Gerontological Nursing, 38(4):18-23.

[8] Rantz M, Scott SD, Miller SJ, Skubic M, Phillips L, Alexander G, Koopman RJ, Musterman K & Back J, "Evaluation of Health Alerts from an Early Illness Warning System in Independent Living," Computers, Informatics, Nursing, 2013, 31(6):274-280.

[9] S Intille, K Larson, E Munguia-Tapia, J Beaudin, P Kaushik, J Nawyn, et al, "Using a live-in laboratory for ubiquitous computing research", Proc. of Pervasive Health, 2006.

[10] Kidd CD, Orr RJ, Abowd GD, Atkeson IA, Essa B, MacIntyre B, et al., "The Aware Home: A living laboratory for ubiquitous computing research", Proc of the 2nd International Workshop on Cooperative Buildings-CoBuild'9, 1999.

[11] KZ Haigh, LM Kiff, G Ho, "Independent Lifestyle Assistant: Lessons learned", Assistive Technology, 2006, pp.87-106.

[12] D Heise, M Skubic, "Monitoring pulse and respiration with a non-invasive hydraulic bed sensor", Proc. 32nd Annual International Conference of the IEEE Engineering in Medicine and Biology Society. Buenos Aires, Argentina. 2010.

[13] R Beckwith, "Designing for ubiquity: The perception of privacy", Pervasive Computing, 2003, pp.40-6.

[14] Galambos C, Skubic M, Wang S & Rantz M, "Management of Dementia and Depression Utilizing In-Home Passive Sensor Data," Gerontechnology, 2013, 11(3):457-468.

[15] Hajihashemi Z, Popescu M, "An Early illness recognition framework using a temporal smith waterman algorithm and NLP", AMIA 2013, Chicago, November, 2013.

[16] Hajihashemi M, Popescu M, "A Multidimensional Time-Series Similarity Measure with Applications to Eldercare Monitoring", IEEE Journal of biomedical and health informatics, to appear, 2015.

[17] Skubic M, Guevara RD & Rantz M, "Automated Health Alerts Using In-Home Sensor Data for Embedded Health Assessment," IEEE Journal of Translational Engineering in Health and Medicine, 2015, 3:1-11.

[18] Popescu M, Mahnot A (2009). Acoustic Fall Detection Using One-Class Classifiers. Proc. 31th Int. Conf. of the IEEE EMBS, Minneapolis, MN, Sept 2-6, 2009, pp. 3505-3508.

[19] Popescu M., Mahnot A., "Early illness recognition in older adults using in-home monitoring sensors and multiple instance learning", IDAMAP-10, Washington DC, Nov 12, 2010, pp. 21-25.

[20] Hajihashemi Z., Popescu M," Detection of Abnormal Sensor Patterns in Eldercare", E-Health and Bioengineering Conference (EHB), 21-23 Nov. 2013, Iasi, Romania, 2013, pp. 1-4.

[21] Bailey T.L., et al., "MEME SUITE: Tools for Motif Discovery and Searching", *Nucleic Acids Research*, 2009, Vol. 38, pp. 202-209.

[22] Z. Hajihashemi, M. Popescu, "A New Illness Recognition Framework Using Frequent Temporal Pattern Mining", UBICOMP 2014, Sept. 15, Seattle, WA, 2014, pp 1241-1245.

[23] Hajihashemi Z., Popescu M., "Early Illness Recognition Using Frequent Motif Discovery", Proc. of IEEE EMBC 2015, Aug. 25-30, Milan, Italy.

[24] Popescu M., Florea E., Linking Clinical Events in Elderly to In-home Monitoring Sensor Data: A Brief Review and a Pilot Study on Predicting Pulse Pressure", Journal of Computing Science and Engineering, vol. 2, no. 1, March 2008, pp. 180-199.

[25] Popescu M., Mahnot A., "Early Illness Recognition in Older Adults Using In-Home Monitoring Sensors and Multiple Instance Learning", Methods of Informatics in Medicine, 2012 Aug 7;51(4):359-67.